Interacting with Temporal Data

at CHI 2009

April 4, 2009
Boston, Massachusetts
Organizers: Wendy Mackay, Max Van Kleek, Aurélien Tabard

Interactively Exploring Time-Oriented Data

Wolfgang Aigner

Silvia Miksch

Alessio Bertone

Alexander Rind

Tim Lammarsch

Dept. of Information and Knowledge Engineering (ike),
Danube University Krems, Austria

{wolfgang.aigner, alessio.bertone, tim.lammarsch,
alexander.rind, silvia.miksch}@donau-uni.ac.at

www.donau-uni.ac.at/ike

CHI 2009, April 4 – 9, 2009, Boston, MA, USA
ACM 978-1-60558-246-7/09/04.

Abstract

Time is an important data dimension with distinct characteristics that is common across many application domains. This demands specialized methods in order to support proper analysis and visualization to explore trends, patterns, and relationships in different kinds of time-oriented data. The human perceptual system is highly sophisticated and specifically suited to spot visual patterns. For this reason, visualization is successfully applied in aiding these tasks and to date a variety of different visualization methods for time-oriented data exist. However, these methods could be improved by accounting for the special characteristics of time. The main aim of our current research is to account for the complex structures of time in visual representations, analysis, and the visualization process. Especially important are interaction methods that aid analysts when dealing with time-oriented data in visualization systems.

Keywords

Information Visualization, Visual Analytics, Interactive Visual Analysis, Time, Time-Oriented Data

ACM Classification Keywords

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

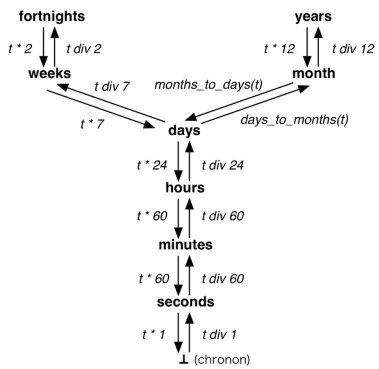

figure 1: Example for some structures of time of the Gregorian calendar.

Introduction

Time is an important data dimension that is common across many application domains, like transport, call centers, retail, production, health care, police, or financial services. Exploring trends, patterns, and relationships are particularly important tasks when dealing with time-oriented data and information. Visualization has been applied to present, explore, and analyze such kind of data for a long time and early representations trace back to the 11^{th} century [10].

In contrast to other quantitative data dimensions that are usually "flat", time has inherent semantic structures, which increase its complexity dramatically. Especially the hierarchical structure of granularities in time, as for example minutes, hours, days, weeks, months, is unlike most other quantitative dimensions. Specifically, time comprises different forms of divisions (e.g., 60 minutes resemble one hour while 24 hours resemble one day) and granularities are combined to form calendar systems (e.g., Gregorian, Business, or Academic calendars). Moreover, time contains natural cycles and re-occurrences, as for example seasons, but also social (often irregular) cycles, like holidays or school breaks. Therefore, *time-oriented data need to be treated differently from other kinds of data and demand appropriate interaction, visual and analytical methods to analyze them.*

To tackle these issues, our research work focuses on the following areas:

- Modeling time and time-oriented data
- Visualization of time-oriented data (in particular the integration of the structure of time and the representation of temporal uncertainties)
- Visualization process of time-oriented data
- Interaction methods for time-oriented data
- Integration of analytical and interactive visual methods that take the structures of time into account
- User-centered development and evaluation of Visual Analytics methods

In the following, we will give a short overview of our past, current, and future research activities in the field of interactive visualization of time-oriented data.

Past Research

CareVis—Integrated visualization of computerized medical treatment plans and patient data

Visualization support for patient data analysis is mostly limited to the representation of directly measured data. Contextual information on performed treatment steps is an important source to find reasons and explanations for certain phenomena in the measured patient data, but is mostly spared out in the analysis process. By the development of *CareVis*, we aimed to fill this gap via integrating classical data visualization and visualization of treatment information [4]. The tightly coupled views use visualization methods well-known to domain experts and are supported by task-specific interaction methods.

PlanningLines—Novel glyphs for representing temporal uncertainties

Planning future activities is a task that we have to face constantly. Since the future is inherently connected with possible uncertainties, delays, and the unforeseen we have learned to deal with this circumstances in everyday life. However, support for temporal indeterminacies is not very well integrated in current

figure 2: Gravi++—Interactive exploration of high-dimensional temporal data [6].

methods, techniques, and tools. To represent and visualize temporal uncertainty concerning the starting and ending times and the duration of actions or events, we have developed novel glyphs, called *PlanningLines* to support project managers in their difficult planning and controlling tasks [5].

Gravi++—Interactive exploration of high-dimensional temporal data
Psychotherapists face the challenge of large amounts of survey data gathered over the course of time in cognitive behavioral therapy. In order to support their task of finding predictors for likely future therapy outcome of anorectic girls, the interactive visualization method *Gravi++* has been developed that utilizes a highly interactive interface for data exploration [6].

Current and Future Research

Systematic view on interactive visualization of time-oriented data
It is enormously difficult to consider all aspects involved when visualizing time-oriented data. Time itself has many theoretical and practical aspects. For instance, time points and time intervals use different sets of temporal relations. Only if the characteristics of the data are taken into account it is possible to generate expressive visual representations and suited interaction methods. We developed a systematic view on the visualization of time-oriented data along four main strands [3]:

- Time: What characteristics of time are considered?
- Data: What is analyzed?
- Representation: How is it represented?
- Task: Which user tasks are supported?

Visual Analytics--Integrating the outstanding visual capabilities of humans and the enormous processing power of computers
Capabilities to both generate and collect data have seen an explosive growth. The need for new methods and tools, which can intelligently and (semi-)automatically transform data into information and furthermore, synthesize knowledge are a core area of the emerging field of *Visual Analytics* [9]. To harness the potential power of a Visual Analytics approach to time and time-oriented data we proposed a conceptual framework that incorporates both, visualization & interaction as well as analytical & mining components [1][2].

In our current research project *DisCō*[1], we aim to develop novel Visual Analytics methods to visually as well as computationally analyze multivariate, time-oriented data and information to discover new and unexpected trends, patterns, and relationships. The main goals of the intertwined visual and analytical methods are to ensure high usability and good control of the integrated mining techniques by applying intuitive visualizations and visual interfaces.

Interaction and the Role of the User
User interactions are one of the most important elements in visualization or even the "heart" as Spence stated [8]. User interaction is even more important in Visual Analytics, as studies, like the one by Saraiya et al. [7] showed: users preferred inferior visualizations with interaction over superior static visualizations. Furthermore, visual representations provide only an initial direction to the data and its meaning, but

[1] www.donau-uni.ac.at/disco; last accessed: Dec 18, 2008

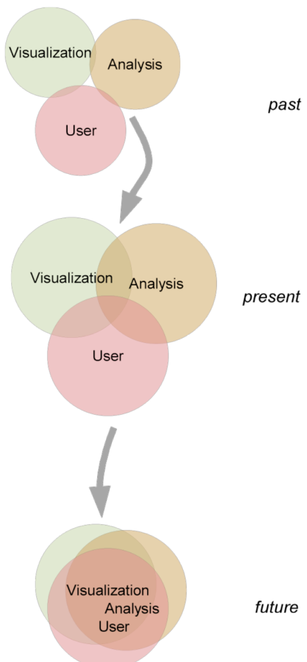

figure 3: To further advance a visually driven analysis of time-oriented data, it is necessary to integrate visual, analytical, and user-centered methods more tightly [1].

through the combination of visual representations and appropriate interaction mechanisms, the users achieve insights into the data [7]. Moreover, it is important that these methods are designed according to users' demands. Interacting directly with the visual representation and the analytical & mining methods provides more control and tighter feedback for the human analyst. This must also include interactive parameterization of both, visual and analytical methods. Navigation methods for large information spaces are decisive for analysis environments that support exploration. In parallel, they should allow for visual overviews as well as the ability to drill down into areas of interest while preserving orientation within the information space. Moreover, user's tasks and goals determine the adequate choice of visualization methods. For example, if we want to identify cycles in the data, suitable representations that reinforce the visual detection of periodic behavior need to be chosen.

In our current research project *VisuExplore*[2], we focus on the aspect of interaction to support medical personnel in their patient data analysis tasks. The project�続s objective is the development of a flexible, interactive visualization environment for time-oriented, medical data and information.

Conclusion

It is widely acknowledged that time is a unique data dimension with distinct characteristics. Many interactive systems deal with time-related aspects. However, these systems could be improved by considering the specifics of time and their implications in a broader sense.

[2] www.donau-uni.ac.at/VisuExplore; last accessed: Dec 18, 2008

References

[1] W. Aigner, S. Miksch, H. Schumann, and C. Tominski. Visual Methods for Analyzing Time-Oriented Data, IEEE *TVCG*, 14, 1 (2008), 47-60.

[2] W. Aigner, A. Bertone, S. Miksch, H. Schumann, and C. Tominski. Towards a Conceptual Framework for Visual Analytics of Time and Time-Oriented Data. Invited paper at *Winter Simulation Conf.* (2007).

[3] W. Aigner, S. Miksch, W. Müller, H. Schumann, and C. Tominski. Visualizing Time-Oriented Data - A Systematic View. *Computers & Graphics*, 31,3 (2007), 401–409.

[4] W. Aigner and S. Miksch. CareVis: Integrated Visualization of Computerized Protocols and Temporal Patient Data. *Artificial Intelligence in Medicine (AIIM)*, 37, 3 (2006), 203–218.

[5] W. Aigner, S. Miksch, B. Thurnher, and S. Biffl. PlanningLines: Novel Glyphs for Representing Temporal Uncertainties and their Evaluation. In *Proc. IV05*, IEEE Computer Society Press (2005), 457–463.

[6] K. Hinum, S. Miksch, W.Aigner, S. Ohmann, C. Popow, M. Pohl and M. Rester. Gravi++: Interactive Information Visualization to Explore Highly Structured Temporal Data, *Universal Computer Science*, 11 (2005), 1792-1805.

[7] P. Saraiya, C. North, V. Lam and K.A. Duca, An Insight-Based Longitudinal Study of Visual Analytics, IEEE *TVCG*, 12, 6 (2006), 1511-1522.

[8] R. Spence, *Information Visualization*, 2nd Ed., Pearson Education Limited, Essex, UK, 2007.

[9] J. J. Thomas and K. A. Cook. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE Computer Science, 2005.

[10] E. R. Tufte. *The Visual Display of Quantitative Information*. Graphics Press, Cheshire, CT, 1983.

Visualizing Temporal Data for Diagnostic Feedback in Adaptive Logistics Systems

Edward Benson
MIT Computer Science and
Artificial Intelligence Laboratory
32 Vassar Street
Cambridge, MA 02139
eob@csail.mit.edu

CHI 2009, April 4 – 9, 2009, Boston, MA, USA
ACM 978-1-60558-246-7/09/04.

Abstract
Adaptive logistics systems pose a difficult challenge for the user interface designer: data is often too complex to show in its raw form and too voluminous to show in its entirety. This paper describes three interactive data visualizations created for a global-scale logistics system and anecdotal lessons learned from its use.

Keywords
Temporal data, planning systems, data visualization

ACM Classification Keywords
H.5.2. Information interfaces and presentation.

Introduction
Adaptive logistics systems manage distributed supply, consumption, and activity plans and dynamically adjust these plans based on real-time feedback. These types of systems help manage everything from retail supply chains to transportation networks, and they are particularly adept in situations where the "plan" might change rapidly, such as disaster relief and military conflict.

These systems pose a difficult challenge for the user interface designer because the data within them is often too complex to show in its raw form and too voluminous to show in its entirety. Projections of future supply chain state are constantly being re-calculated based on feedback from the world (*actuals*), creating new projections and expiring old ones. Data is thus splayed across several alternate timelines of possibility and actuality – projections, past projections, and actuals. In addition, data is multiresolutional: the same basic inventory measurements can be applied to supply truck and continent-wide transportation fleet alike.

Despite the complexity of the data within, the situations where adaptive logistics systems are needed most also have some of the most demanding requirements on the simplicity of user interface. At run time, the interface must be "glanceable," conveying key information to in-field operators in little time; the interface should make problems easy to diagnose; and whenever possible data display should enable exploration without requiring it.

This paper presents three interactive visualizations developed for a global-scale adaptive logistics system. These visualizations attempt to convey information in a simple manner while still preserving the rich temporal and multiresolutional dimensions of the data within. The first two visualizations were implemented atop the Cougaar Multi-Agent Architecture [1], and the third was designed but never implemented. At the end of the paper is an anecdotal list of lessons learned about conveying data in this form.

Three Displays for Diagnostic Feedback

These three displays attempt to help an in-field operator identify and diagnose supply chain problems in a Cougaar-managed logistics operation. Each display is designed to answer a specific question related to this task that helps gather enough information to detect and diagnose the problem.

1. *How long do I have?* Assuming no further replenishment activity occurs, how much longer could I continue following the plan?

2. *Where's my stuff?* What does the supply chain look like, and who has the materials I need?

3. *Should I Trust the Computer?* How accurate have the system's predictions been so far?

To help explain these three displays, consider a fictional scenario during the Hurricane Katrina relief effort of 2005. A coordinator is monitoring the progress of relief operations across a citywide team and notices that the system anticipates problems. She uses Display 1 to drill down and learn that the 9th Ward is now projected to run out of medical supplies. She then uses Display 2 to identify the cause of the supply-chain problem, and finally uses Display 3 to check the running accuracy of the predictions thus far.

How Long do I Have?

A question of importance to anyone dealing with scarce resources over a plan of action is how much longer can they continue to operate if cut off from resupply. We call this *operable time*. When a logistics system contains a fixed timeline, an additional metric called *operational completeness*. An entity is operationally complete at time t if its operable time at t equals or exceeds the remaining plan time.

The display in Figure 1 attempts to answer the question of "*How long do I have?*" by displaying a combination of

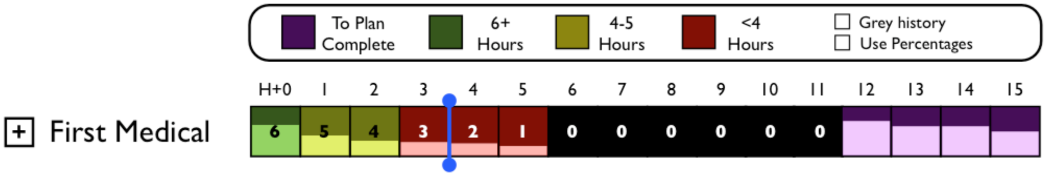

Figure 1. Operable Time for the First Medical Unit

inventory level, operable time, and operational completeness across the duration of the plan. The blue circle-capped bar at the start of hour 4 in the plan represents the current time; all data to the left is actual data from the past, and all data to the right is projected. Each cell contains the operable time for that hour, actual or projected, and is shaded in the background to represent the percent fill of the inventory container. Aggregation is performed by selecting the worst value among the entities being aggregated, although various other aggregation styles can be selected. Colors provide a visual reflection of the operable time, and operationally complete items are shaded a special color (purple) and do not show an operable time.

This display can be expanded hierarchically to reveal increasingly finer-grained components based on interest level. In typical operation, this occurs when a problem is spotted at the aggregate level, and an operator wants to track its cause down to the inventory level.

For example, Figure 1 shows that the First Medical unit is beginning to experience severe inventory problems. Unfolding the First Medical group in Figure 2, we see that the problem is occurring in the 9th Ward. Further unfolding the 9th Ward in Figure 3, we see that the problem is related to medical supplies, not fuel. Further unfolding would reveal more detailed information if desired. This manner of exploration allows an operator to detect an error at just a glance, but then increase the complexity of the display as needed to determine the cause of the error.

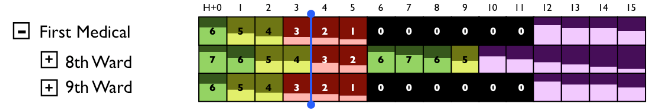

Figure 2. Unfolding the First Medical Unit

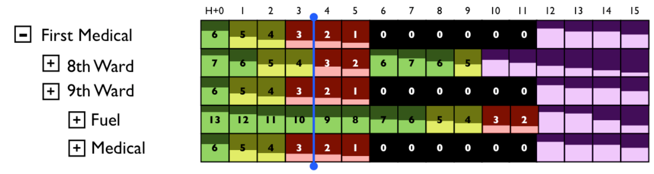

Figure 3. Unfolding the 9th Ward

Where's My Stuff?

After identifying an inventory shortfall, the next task is to identify what happened to the expected supplies. The "Where's My Stuff" chart answers this question by providing a vertical view of a supply chain from the perspective of the consumer experiencing shortfall. Figure 4 depicts movements of a resource throughout a supply chain over time. The blue bar continues to separate the actual past from the projected future. In the operational version of this display, an overhead panel permits the selection of consumer, resource type, and supply channel, as well as several optional parameters that tweak the information being displayed.

The height of the chart represents the total supply of a single inventory item across one vertical supply channel. Supply starts at the top of the chart and flows to the bottom of the chart. The state of "having been consumed" is the final stop in this supply chain and is added to the bottom in white. Two-toned shading is used to denote the difference between supply allocated by the planning system for future transfer down the channel and unallocated resources. For the purposes of this paper, resources allocated elsewhere are not considered. Icons decorate the chart to represent

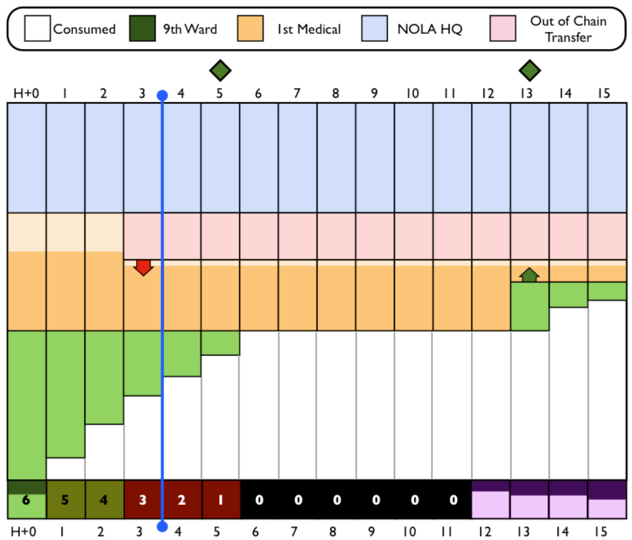

Figure 4. Where's My Stuff?

replenishment activities: diamonds at the top for *planned* inventory transfers, green arrows for *actual or expected* inventory transfers, and red arrows for *unexpected, out of channel* transfers. At the bottom of the chart the operable timeline is re-displayed for the selected entity and inventory item.

Looking at the chart in Figure 4, we see that an unplanned inventory transfer occurred at Hour 3, causing a drop in the overall supply of the 1st Medical supplier. At hour 5, 2 hours in the future, we see that a planned replenishment is not expected to occur. Without leaving this chart to explore more, one can guess the problem and suggest a solution. It appears that 1st Medical, the immediate supplier of the 9th Ward, has been diverted unexpectedly to replenish another location and will be unavailable for the planned operation at hour 5. This is likely the logistic cause of the expected medical inventory shortfall seen in Figure 3. A potential fix would be to contact NOLA HQ, which appears to contain a full inventory of uncommitted supplies, and ask for a direct inventory transfer.

Should I Trust the Computer?

The final step that our hypothetical operator needs to make before contacting NOLA HQ is to check the accuracy of the system's past projections to know how much trust to place in its current projections. We define the *forecast horizon* as the distance into the future that a system is measured to have made accurate predictions. Notions of accuracy, of course, depend on the user and situation, so the display shown here uses a customizable color quantization to convey accuracy.

Figure 5 shows the forecast horizon chart, which helps an operator determine the forecast performance for a particular inventory item. Present time, indicated again by the circle-capped bar, has been moved forward to Hour 8 to provide a better diagram for explanation. The operable timeline for the selected consumer and inventory pair marks the horizontal axis, and the vertical axis represents relative time in the past.

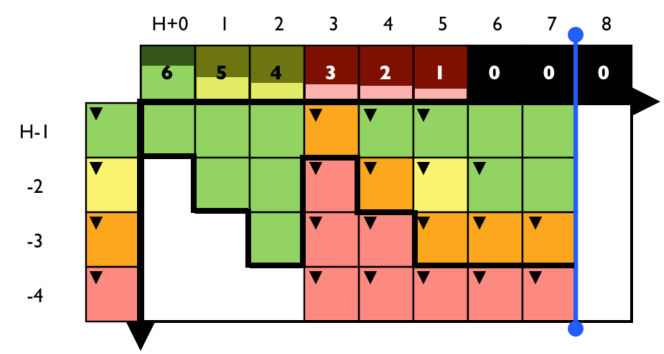

Figure 5. Forecast Horizon

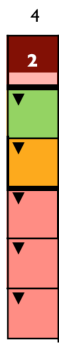

Figure 6. Forecast Horizon at Hour 4.

The colored shading of each cell represents the quantized measure of how acceptable the deviation between projected and actual inventory levels is at that hour. At Hour 4, for example, we see that the projection one hour prior was good (green), at two hours prior was mediocre (orange), and at 3+ hours prior was bad (red). An additional marking is added if the actual data showed *less* inventory than projected.

From this chart, the operator can see that hour 3 caused serious forecast errors but the system seems to have leveled out to a forecast horizon of 3 hours afterwards. Barring another unexpected event, it is then reasonable to trust the current projections at 3 hours into the future but no further.

Observations and Lessons Learned

In iterating over the first two visualizations with users of the system and planning for the third, a number of anecdotal observations arose that are worth repeating.

Target Visualizations to Your User's Stories
General-purpose displays are not always beneficial in time-critical, information-rich situations. The many possible ways that complex temporal data can be displayed is exciting to a computer scientist but a potential nightmare for an in-field operator. We found that some of the most well-received visualizations were those that responded directly to a story a user told us about a time they needed information but could not figure out how to retrieve it from the system.

Lines Plots Considered Dangerous
When dealing with temporal data collected at discrete points in time, beware of using line charts. The interpolated line drawn between points implies information that might not be supported by the data. It is *especially* erroneous with regard to supply-chain transfers.

Support Pay-As-You-Go Complexity
We found it useful for each display to start with the simplest possible variant for the logged-in user's role. The operable time chart for a system-wide planner, for example, would start entirely collapsed. This enables the operator to quickly glance at the display to determine *if* further inspection is necessary, and only then drill-down into more detailed specifics.

Users Want Hyperlinks
The more clickable areas on a chart, the happier our users were. When visualizations are targeted to answer specific question, elements on those visualizations can link to the obvious "next steps" that the data implies.

Color is Powerful; Color is Distracting
We found that quantizing key information into a color shading a great way to create glanceable displays, but the multi-colored displays can be overly-confusing. Users appreciated the ability to turn all data in the past into grayscale, or even remove it completely.

Acknowledgements

This paper describes work performed while at BBN Technologies. The author extends great thanks to Jeffrey Berliner, who played a critical role in devising and refining these visualization schemes.

Citations

[1] A. Helsinger, M. Thome, T. Wright, "Cougaar: A Scalable, Distributed Multi-Agent Architecture", IEEE Conference on Systems, Man and Cybernetics. 2004.

From Multimedia to Information Management: Using Semantic Time in Interaction Design

Jan Borchers
Media Computing Group
RWTH Aachen University
52062 Aachen, Germany
borchers@cs.rwth-aachen.de

Thorsten Karrer
Media Computing Group
RWTH Aachen University
52062 Aachen, Germany
karrer@cs.rwth-aachen.de

Jonathan Diehl
Media Computing Group
RWTH Aachen University
52062 Aachen, Germany
diehl@cs.rwth-aachen.de

CHI 2009, April 4 – 9, 2009, Boston, MA, USA
ACM 978-1-60558-246-7/09/04.

Introduction

We like to think that we have a solid understanding of time—we can plan, schedule, and synchronize anything through rigid temporal structures and a universal syntax of time. Our conceptual model, though, rarely follows the rules of this artificial construct. We rather think (and speak) in terms of *semantic time*, framing events in meaning instead of numbers. "Having a meeting right after lunch", "falling asleep during the second movement of Beethoven's 9th", or "pausing a video at the point where the player commits a foul" are examples of how we use semantic time every day.

On the following pages, we will outline how we employ the concept of semantic time in different aspects of our research. First, we discuss interacting with time-based media such as audio and video without relying on a fixed time base: we present the *Semantic Time Framework* and *PhaVoRIT* as toolkits and *Personal Orchestra* and *DRAGON* as example applications. Then we discuss how humans consider time, and relate it to the area of information management where data is traditionally bound to precise syntactic time.

Time-based Media

A large part of the media that are created, shared, and consumed today are time-based media like videos, music, or human speech. Their contents only unfold linearly over time, and a single frame or sample has little meaning by itself. Established techniques to navigate, search, and play time-based media through syntactic time often do not directly support the users' tasks. Finding the scene where the Death Star blows up or listening to a sonata without the repetitions are tasks that call for a way to express time that regards the semantics of the medium itself.

Developing semantic time representations for multimedia and creating interaction techniques that support them is a challenge we would like to discuss in this workshop. Solving the problems in this domain could drastically change the way we interact with time-based media.

In order to investigate how semantic time can improve the way we interact with these types of media, we have approached the problem on multiple layers:

- On the implementation-centric layer of APIs and toolkits, we created the *Semantic Time Framework (STF)* [4]. It simplifies developing interactive multimedia applications that allow users to interact with media using semantic time. STF abstracts different media sources as *streams,* each having a different sample density and flow rate. Semantically meaningful points in these streams carry special markers, thus making it possible to quickly navigate to any semantic point in time or to easily control their rate of succession. Semantic time information in one stream can also control how another stream is processed. This allows for intelligent synchronization between different time-based media, e.g., synchronizing the beats of a musical piece that is being played back to the gestures of a conductor.
- In order to control the playback rate of a wide range of time-based media, we developed *PhaVoRIT* [3], an advanced phase vocoder based audio time-stretching engine. It facilitates arbitrarily changing the speed of audio data without altering the pitch, while hardly degrading the quality of the signal. The operation takes effect with minimal latency, allowing it to be used in interactive applications.
- *Personal Orchestra (PO)* [1] is an example of how an application can make use of the lower layer support for semantic time to present new interaction possibilities to the user. Building upon STF and PhaVoRIT, PO captures a user's conducting gestures to control the playback of an audio/video recording of the Vienna Philharmonic Orchestra. The interaction is based on semantic time, letting the user determine the progression of the beats and bars of the music as she desires.
- A different approach to semantic time on the application layer is *DRAGON* [2], our direct manipulation video navigation technique. Browsing through videos is usually done using the timeline slider, an instrument to control the syntactic time of the medium. The goal of the user, however, is often better expressed in terms of the plot or a certain constellation of objects in the scene. Examples include the specific frame where a sprinter crosses the finish line, or where a car passes a traffic light. By directly dragging objects in the scene along their movement trajectory, *DRAGON* enables users to quickly and precisely navigate to a specific point in the video timeline where an object of interest is at a desired location.

Information Management

Personal Information Management (PIM) is an umbrella term for the acquisition, organization, and retrieval of all information that is relevant for our everyday lives. Time is an essential aspect of PIM but the way it is handled by today's systems is different from how we as humans think about time. To improve usability, we must overcome this mismatch and create a more natural representation of time. Consequently, we need to understand how humans think about time. Current literature mostly concentrates on time perception, so we have come up with our own set of strategies that humans employ when discussing time:

1. Syntactic: time is defined on a fixed temporal scale (2:49pm). This basic form of time is primarily used for synchronization with external factors or people.
2. Imprecise: time is approximated ("We will be done in a few minutes"). Approximation is used for defining times ("I need the report today.") and for informing about time ("The train arrives shortly before 3.") in a casual way.
3. Anticipatory: time affords foresight ("I need to get up early enough to arrive at the airport in time") or hindsight ("I should start handing back the corrected exams soon"). The given event might require time for pre and/or post processing, which must be considered.
4. Relative: time is related to other time (5 minutes after…). Instead of using a specific time, events are related to one another ("Let's meet 5 minutes after closing time").
5. Ordered: time is put in order ("First add sugar, second add milk, …"). Ordered time, as opposed to relative time, cannot be translated to a specific point in time because it relies on the duration of the previous event and might not occur immediately.
6. Contextual: time is related to other events ("I need to ask Peter about this the next time I see him."). For contextual time, it is unclear when or if it will happen.
7. Abstract: time is abstracted to greater happenings ("The music of the 60s was great!"). Abstract time only considers an event in the context of a greater happening with no reference to a specific point in time.
8. Periodic: time occurs repeatedly ("We meet every Monday."). Periodicity is defined through rules and exceptions ("I like to go swimming every other day except on the weekends.").
9. Combined: strategies may be combined arbitrarily ("I go running every Tuesday and Thursday after work, unless it is raining.").

This list can only be considered as preliminary and should be backed up by concrete research, but we find it useful for starting a discussion about how computer interfaces should consider time. Some of these strategies are already reflected in existing systems:

- Imprecise: Apple displays the remaining time for system tasks, e.g. copying files in Finder, imprecisely.
- Anticipatory: Calendars allow the user to define reminders, which are triggered before the task occurs.
- Relative: When entering dates on the web, often a calendar is displayed, allowing the user to select a day relative to the one marked as today.
- Ordered: To-do applications commonly allow prioritization of tasks, which puts them in order.

- Contextual: Some innovative task management applications allow defining a context for tasks, which can be used to show the task at the appropriate time.
- Periodic: Calendar applications usually allow defining simple rules for repeating events.

The support of these strategies seems uncoordinated, though: (1) it is not universal: at some point in the interface a strategy is supported while at a different point it is not; and (2) strategies cannot be combined: The user may choose one form or the other but not an interdependent combination of both.

Discussion

By moving away from rigid time structures and supporting humans in their natural way of dealing with time, we hope to improve the usability of interactive systems.

An important challenge here is to create a framework that defines how to represent time in interaction design, for which we hope to contribute in the areas of interacting with time-based media and information management. This framework should aim at unifying the experience of working with time in interactive systems in a natural way by giving guidance to developers and designers. To validate this framework, we intend to extend our examples with new interaction metaphors and compare them to existing systems.

Our goals for the workshop are:

- to share our ideas and results surrounding the concept of semantic time.
- to identify new natural ways of interacting with time-based media.
- to get critical feedback about the human strategies for dealing with time.
- to design example representations of time that reflect the human strategies.

During the workshop, we offer to share our insights from the development and integration of toolkits, and the design and evaluation of new interaction techniques, as well as our experience in deploying multiple demonstration systems, showing new interactive technology in realistic situations.

We are very interested in actively participating in a community of researchers interested in temporal interaction design and visualization, as we are in working towards a special issue for a journal on this topic.

References

[1] Borchers, J., Lee, E., Samminger, W., Mühlhäuser, M. Personal Orchestra: A real-time audio/video system for interactive conducting. *ACM Multimedia Systems Journal Special Issue on Multimedia Software Engineering*, **9**(5):458-465, March 2004.

[2] Karrer, T., Weiss, M., Lee, E., and Borchers, J. *DRAGON*: a direct manipulation interface for frame-accurate in-scene video navigation. *Proceedings of CHI '08*, 2008, 247-250.

[3] Karrer, T., Lee, E., Borchers, J. PhaVoRIT: A Phase Vocoder for Real-Time Interactive Time-Stretching. *Proceedings of the ICMC '06*, New Orleans, USA, November 2006.

[4] Eric Lee. *Semantic Time Framework.* PhD thesis, RWTH Aachen University, 2007

Using Temporal Patterns (T-Patterns) to characterize routine tasks

Oliver Brdiczka
Palo Alto Research Center (PARC)
3333 Coyote Hill Road
Palo Alto, CA, 94304, USA
oliver@parc.com

Norman Makoto Su
Department of Informatics
University of California, Irvine
normsu@ics.uci.edu

Bo Begole
Palo Alto Research Center (PARC)
3333 Coyote Hill Road
Palo Alto, CA, 94304, USA
bo@parc.com

Interacting with temporal data workshop at CHI 2009.
http://temporal.csail.mit.edu/

Abstract

We describe the use of a statistical technique called T-pattern analysis to derive and characterize the routineness of tasks. T-patterns provide significant advantages over traditional sequence analyses by incorporating time. A T-pattern is characterized by a significant time window (critical interval) that describes the duration of this pattern. Our analysis is based on data collected from shadowing 10 knowledge workers over a total of 30 entire work days. We report on the statistics of detected T-patterns and derived correlations with participant perceptions of workload, autonomy, and productivity.

Keywords

Temporal patterns, T-patterns, routine tasks

ACM Classification Keywords

H.5.3 Group and Organizational Interfaces: Theory and models, evaluation/methodology, asynchronous interaction, synchronous interaction

Introduction

We propose the utilization of a statistical technique called T-pattern analysis to derive and characterize the routineness of tasks. Based on the assumption that constant duration of steps or events is distinctive for a routine tasks, the T-pattern analysis isolates patterns that are significant in their temporal configuration. Applied on observational data from application and document usage as well as email communication of 10 employees, we could show the number and length of detected T-patterns correlate with perceived workload, autonomy, and productivity.

Uncovering Temporal Patterns

One of our goals in this paper is to advocate the usage of T-patterns in the analysis of human-computer interaction.

We assume that routine tasks can be characterized by specific recurrent actions that are executed within nearly constant time intervals. In order to detect such patterns, we used a probabilistic temporal pattern detection method, called T-pattern detection [1]. T-patterns are recurrent events that occur within a similar temporal configuration (critical interval, CI). The T-pattern detection algorithm uses a statistical test (CI test) that reveals whether the temporal distances between all occurrences of two events are random or not (with respect to a specified p-value). The CI test is based on the null hypothesis that two events A and B are independently and purely randomly (Poisson) distributed over the observation period. The test is applied on all observed temporal distances between the two events A and B and their frequencies, identifying the distances that are supposedly not random according to the specified p-value. Beginning with testing all possible pairs of basic events and thus isolating significant basic patterns, the T-pattern algorithm then successively constructs larger patterns by combining events and significant basic patterns that have been found. Thus, the T-pattern detection algorithm identifies highly significant, hierarchically arranged T-patterns that are composed of statistically related events that repeatedly appear in the same, relatively invariant, temporal configuration.

T-pattern parameters

Our implementation of the T-pattern detection algorithm is based on the description by Magnusson [1]. The relevant parameters used to obtain the results described in this section are:

1. Minimum Occurrence = 2—specifies that a given pattern must occur at least twice to be included in the results.

2. Significance Level = 0.05—specifies the probability that a given pattern would occur in a random (Poisson) distribution of the current data set.

3. Maximum pattern length = 4—specifies the maximum number of action or events that a pattern can be composed of, in order to reduce the complexity of the algorithm and to filter only reasonable pattern sizes.

Derived T-pattern statistics

The T-pattern detection algorithm identifies a number of T-patterns (N_t) that are significant per task. In addition, the significant minimal and maximal temporal length (d_1 and d_2) for each T-pattern is reported, that is, if A is an earlier and B a later component of the same recurring T-pattern, then, after an occurrence of A at t, there is an interval [t + d_1; t + d_2] ($d_2 > d_1 > 0$) that contains at least one occurrence of B. Figure 1 shows an example case, reporting on the T-pattern Outlook-Word. The following statistics can be derived:

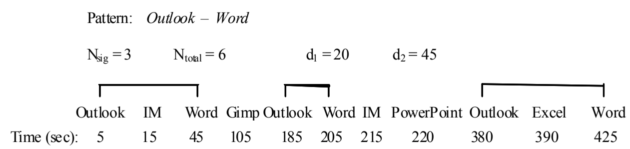

Figure 1. T-pattern analysis detecting a pattern of Word following Outlook from 20 to 45 seconds

1. N_t = number of different T-patterns found per task, which refers to the number of T-pattern classes that have been identified to be significant

2. *minL* = $\frac{\sum_{i=1}^{N_t} d_{1i}}{N_t}$ = average minimum temporal length of the T-patterns per task.

Data Collection

Data was collected in situ, shadowing a total of 10 employees. All but one informant (who was shadowed for only two whole days due to scheduling constraints) was shadowed for three whole work days. The observer would meet the informant upon their arrival to work and follow the informant as closely as possible until the end of the business day. Using a paper notepad, the researcher would label, to the second, user tasks and their start/end times. Parallel to the shadowing, logging software has been installed on each participants PC recording application usage. The application usage that has been recorded included application name, window type and size, documents opened, email sender and recipients. Shadowing notes and PC logging data were merged and synchronized. This resulted in a dataset that contained application usage events from the PC logging data that were annotated with their appropriate task label from the shadowing notes.

At the end of each shadowing session, we administered 3 surveys aimed at measuring informant's stress. Final composite scores of productivity, autonomy, and workload were calculated from these surveys for each informant.

Data Analysis and T-pattern statistics correlations

The T-pattern analysis has been conducted separately for the following fields of the logging data trace:

1. Application and window class (appwclass)
2. Active Document (doc)
3. Email sender – recipients (email)

Application and window class refers to the current application (e.g. EXCEL.EXE) and the current window class (e.g. ConsoleWindowClass). Note that you normally have several window classes per application, providing a higher granularity of observation of the user's actions. The active document refers to the document that has current focus. Email sender and recipient refer to the event when the user receives or selects an email. Sender and recipient ID of this email are the input events of the T-pattern analysis. Table 1 reports on some interesting correlations between the T-pattern statistics and the psychometric stress measures. Significant pairs ($p<0.05$) are in bold, trends ($p<0.10$) are underlined. The corresponding *p*-values are in parenthesis.

		Workload	Autonomy	Productivity
appw class	N_t	0.33 (0.10)	0.07 (0.73)	0.07 (0.72)
	minL	-0.06 (0.75)	-0.15 (0.47)	-0.16 (0.43)
doc	*N*	**0.45 (0.04)**	0.35 (0.12)	0.35 (0.12)
	minL	0.13 (0.58)	0.18 (0.43)	0.12 (0.59)
email	N_t	-0.18 (0.39)	0.08 (0.70)	-0.03 (0.87)
	minL	-0.20 (0.33)	-0.34 (0.10)	**-0.48 (0.02)**

Table 1. Pearson's *r* correlations between psychometric and T-pattern statistics

The number of T-patterns from document usage correlates positively with workload stress. A similar trend can be found for T-patterns from application window classes. The minimal length of T-patterns from email sender/recipients is negatively correlated with productivity. A similar trend exists for autonomy. From this, we might conclude that the longer the minimal length of sender-recipient patterns, the less the informants felt autonomous and productive.

Workshop Objectives

The results presented here are a first step in creating applications that provide users with temporal awareness of their own task performance. These analyses and visualizations are useful to analysts but are likely to be difficult for end users to interpret. Prior research in temporal pattern representations (e.g., [2], [3]) has not generated generalizations of design principles that will apply to the temporal structures of T-patterns. The user-interaction research community needs to map out the requirement space for temporal pattern representation and generate a set of principles that can be applied when end users need to interpret and manipulate temporal structures.

Conclusion

This paper aimed at introducing the T-pattern analysis for detecting temporal patterns in routine tasks. A first correlation analysis of the T-pattern statistics from application and document usage as well as email communication showed that these can be used to indicate workload, autonomy and productivity. We believe that T-patterns are a useful tool for uncovering temporal structure in event-based data. T-patterns unveil a window of constant durations between events, which makes them particularly suitable to characterize routine work. In participating in this workshop, we aim to contribute to the set of requirements for temporal interaction principles and to broaden our understanding of existing temporal interaction principles.

References

[1] M. S. Magnusson. Discovering hidden time patterns in behavior: T-patterns and their detection. *Behavior Research Methods, Instruments, & Computers*, 32(1):93–110, 2000.

[2] James "Bo" Begole, John C. Tang and Rosco Hill, "Rhythm Modeling, Visualizations and Applications," *Proceedings of the 2003 Symposium on User Interface Software and Technology* (UIST 2003), pp. 11-20.

[3] Fisher, D. and Dourish, P. Social and temporal structures in everyday collaboration. *In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI 2004) ACM, New York, NY, 551-558.

Time Machine Computing?

On the use of time and memory in everyday settings

Gaston R. Cangiano and James D. Hollan
Distributed Cognition and Human-Computer Interaction Laboratory
Cognitive Science Department,
University of California at San Diego
{gaston,hollan}@cogsci.ucsd.edu

INTRODUCTION

There is a long history in HCI research and a recent resurgence of interest in recording personal workstation activity as a means to access and mine past behavior. This includes early work on history-enriched digital objects [1], lifestreams [2], and time machine computing [3]. Currently a number of projects are exploring the capture of extensive digital life archives and the visualization of activity-based histories [4, 5]. These projects and a growing interest in personal information management (see [6] for a recent survey) highlight the importance of understanding effective ways to manage and interact with personal histories of information-based activities. Surprisingly, there has been few empirical evaluations of existing systems ([7] and [8] are notable exceptions). We are interested in whether logging technologies can help in exploiting past activities and whether they can be effectively integrated into everyday work practice.

The experimental literature on everyday memory richly documents the myriad ways people rely on memory of the past in everyday life. But in terms of personal computing activity, we know little about how people use records of the past in this increasingly important setting. People do co-opt the environment around their computers to help reinstate past context. Display screen edges, for example, are frequently filled with post-it notes. Desks are organized in ways to help people to resume work in the future. Scientists and other professionals journal recent events in notebooks next to their computers or in files to capture thoughts and ideas that will remind them of their activities. But these and similar facilities only capture a partial views of the rich structure of activities and their contexts and they are typically not well integrated with those activities.

In this paper we highlight two issues we see as central to the workshop theme: (1) a need for careful observation and documentation of people's everyday workstation activities in order to understand how people adapt and make use of current facilities to aid memory and to help reestablish past contexts, and (2) a need to explore new designs to unobtrusively capture activities and more effectively support reestablishment of previous contexts in ways that leverage knowledge of how human memory works.

We first briefly present ethnographic work we have conducted over the past year at a law office to observe and record work practices. We use this data to illustrate how currently available tools lack support for reestablishing earlier contexts. We then briefly review selected literature from experimental psychology on the mental representations of time and discuss its relevance to designing more effective visual representations of time and interaction paradigms.

FIELD WORK

Our field site is an Immigration Law firm located in the San Francisco Bay Area. This medium sized firm has been in operation for over three years. Its members have worked together as a team for almost a decade, starting as employees of a large immigration law firm before leaving to start their current company. Shaped by a history of working together they have well-established work practices for assisting clients with immigration issues. The work flow of the office is smooth and efficient.

Unlike other types of law firms that need to keep track of billable hours, billing in this office is done on a per visa or work permit basis. The focus is thus on keeping track of deadlines, client communication, and documentation. It is important for paralegals to keep case logs, not with the time spend on each case, but rather with a timeline of chronological interactions with clients, filing deadlines and follow up alerts. Because of the emphasis on "time-lining" rather than "time-tracking", both paralegals and attorney often multitask more frequently than they might if specific time was being charged to each client. They can address issues of many different clients in an interleaved fashion, not having to worry how much time they spend on each email or phone call. A common pattern that emerges from this interaction is that everyone usually works on multiple cases at the same time; the only exception is when a case is being compiled just prior to a filing deadline or when a new case is opened and substantial data entry is needed.

A Work Scenario

SH is a senior attorney at the Law Firm. The role of attorneys in this office is to address legal questions from the paralegals and clients, review cases and letters, and to keep current with legal research, the industry, and government legislation. A significant portion of SH's daily activity consists of responding and sending emails to clients. In addition, she spends a substantial amount of time sending and reading instant messages (IM) with her office colleagues regarding cases. Inter-office communication via IM is an integral part of the work setting, even when participants are within ready conversation distance from each other.

Answering a client email typically requires SH to place the case in perspective, that is, to refresh her memory of the history of the case. Reviewing a case's history involves retrieving previous correspondences, retrieving information from a client database and consulting with the other attorneys and paralegals involved. The latter is usually accomplished via IM. Frequently a "reply" window remains open in the background for extended periods, while she collects and refreshes her understanding of the information about the case needed to draft a reply. Consider the following excerpt taken from an interview in which she was seeing a history of recordings of her display screen:

> `this is a client writing with a question and I am responding to the client but I think what am I doing let's see I probably don't know all the information usually before I have to respond to a client I have to contact via IM another legal assistant which is what I seem to be doing`

Imagine that a month later she needs to recall the reasons she gave a particular piece of legal advice that morning. This is a common scenario in the Law Office because cases share similarities and therefore previous resolutions can be applied to new cases as part of providing legal advice or drafting petition letters. Her original advice developed after reading a set of earlier emails and discussing related issues with others in the office. The only records available to assist the process are the resolutions of the case, noted in a database, and milestone comments annotated along the way. Perhaps a very meticulous paralegal would record additional comments in the database, but the reality is that with the high load of cases at a successful law firm, this is seldom done. The nuances of each episode remain only in SH's memory, everything else is lost.

An important design challenge therefore is how to use digital activity recordings to assist her in quickly reconstructing the reasoning behind her previous legal advice. Currently, it takes considerable effort to bring the separate pieces of information together in order to reassemble the context of a past episode. The situation is exacerbated by the tools available today which provide virtually no support for this type of behavior. The increasing complexity of work and frequent use of multiple independently designed applications limit the ability of developers of any single application to provide comprehensive support for real-world practices. This results in increased cognitive load and fragmented contexts. Human memory is finely tuned to reconstruct the past but this fragmentation removes access to the perceptual cues associated with the temporal unfolding of the original episode that might help trigger memory.

The situation is illustrated in the following interview excerpt (emphasis added):

> `… we are shuffling about in general for each of us about 100 to 150 cases and there is no way we remember so ... we're ... the ... one database that we have online database does not allow window more than one window to be open ... so we have one part of the database that has a log of all the conversations we had all the emails exchanges strategy notes ah ... and then we have the outlook calendar or task list for updating what we need to do next and the case log ... the online database hmmm ... it takes too long and possibly you constantly have to type two different fields enter first name second name pieces of it to get that person up and then once that screen is up there are many different ... chambers within that screen to ... just to find out what had occurred ... let's say a month ago ... so if there is an` **`issue`** `in the case it's buried in that database ...`

The "issue" mentioned above consists of the very nuances and details about a case which most likely are not explicitly annotated, but rather need to be reconstructed by assembling the separate pieces of a case history. These pieces likely consist of emails, conversations, and annotations which are currently dispersed across several different applications and places within a database.

We argue that in order to design tools to support the kind of behavior necessary for reinstating past contexts, it is important to understand the cognitive processes involved in evoking the meaningful structure of past activity. The visual cues and temporal structure available to us during the performance of an activity guide our attention to what we will remember. This kind of information must be made available again to some degree when trying to access the past. Activity images might provide the perceptual cues needed to trigger memory through the process of *recognition*. Activity recordings might also provide the temporal structure of entire episodes. One can think of *reinstating* context as an act of "reperceiving", that is, perceiving *again* the same visual and temporal stimuli used to guide attention and memory during the original performance of the activity.

Activity recordings to be effective therefore must satisfy two important cognitive constraints: first, they must cut across the artificial boundaries established by applications to provide a unified temporal flow from the available fragmented contexts; secondly, they must engage episodic memory through the use of visual and temporal cues in order to engage the cognitive mechanisms responsible for reinstating the mental context of previous activity.

ON MEMORY AND REPRESENTATIONS OF TIME

The literature on autobiographical memory indicates that people excel at remembering histories of events [8, 9, 10, 11]. People can easily reconstruct information about chronologies and the ordering of personal events that are meaningful in their lives [8]. This is referred to as preserving the *temporal* flow of events [9]. At the same time, we seem to have fairly poor memory for exact dates and locations [8,10]. This is not to say that we cannot remember dates, but that we are better at recalling the past based on *thematic* cues rather than calendar information. The organization of autobiographical memory is structured around personal *themes* and these in turn constitute the retrieval cues used to access the past.

Experimental work has shown that the classic exponential memory decay curve becomes linear when memory for the past is cued based on familiar items, such as faces or personal names, rather than locations or dates [8]. Memory over long periods of time is more stable if cueing is done based on familiar categories (in the field of HCI, Sellen et al found a similar effect [7]). These effects are robust, and produce consistent biases when people try to estimate the time of past events [12]. This has led theorists to conjecture that mental representations of time are not isomorphic to calendar time or even single timelines, but rather, multiple-scaled multiple-entry representations [11]. In other words, we use a multiplicity of cues in our memory to *reconstruct* the time of past events. The past is therefore represented at different timescales and via different thematic categories. These entry points can be considered memory “landmarks”, and play the role of structuring the organization of memory.

The idea of memory “landmarks” is widespread in the psychological literature. Landmarks have been shown to produce significant effects on accuracy and vividness of recall in experimental settings [10, 13, 14]. But what exactly is a landmark and how is it selected? This question is difficult to answer, because the human memory system cannot anticipate what events will be useful in the future to recall the past. The most obvious candidates are dramatic public events, known as “flashbulb” memories. But personally meaningful events are also salient landmarks. These events however have clear limitations for the purposes of designing software tools to access the past [15]. Their main drawback is that they require users to explicitly tag events as meaningful, something we don’t naturally do in our normal activity and which would include additional cognitive load.

Another important type of landmark to consider is cyclical events. At a finer grain of temporal resolution, recurring events are also used as landmarks, not because of their saliency but because of their *stability*. For example, the fact that one knows that on Wednesdays there is staff meeting at work, or that on Fridays one usually telecommutes, can be used as memory retrieval cues for reconstruction of the past. The word *reconstruction* here takes on a distinct meaning. The literatures on autobiographical and episodic memory both emphasize the distinction between the processes of re-experiencing the past and those of recalling it. Over longer periods of time the process of reconstructing the past becomes increasingly significant [20, 10].

Cyclical landmarks may play an important role in the process of context reconstruction, but in order to understand how we can use them for design, it is important to investigate ways to support *navigating* the past. Some interesting research projects dealing with personal information such as MyLifeBits [3] and Feldspar [17] have explored new concepts like “pivoting” or “orienteering” as a means to explore and search the past. Both approaches view searching as a combination of search and browse. Pivoting for instance, allows users to start a query using meta-tags, and then use the results to narrow search to smaller subsets of data. The intermediate results from a query can be browsed so they serve as a memory aids. Cyclical landmarks can further enhance this behavior by allowing users to first orient themselves by using familiar cues from their everyday activities. The process of getting back into the past can then perceived as a mixture of using what we know about ourselves to initiate retrievals and then reinstating the past by browsing and replaying past activity. This behavior shares much in common to what we do when reconstructing the past using episodic memory.

Timelines are a common representation of time. While explored primarily as a means to support *reflection* about the past [4], they potentially can also serve context reinstatement purposes. We conjecture that it will be important to explore timelines that are not isomorphic to calendar time, but rather reflect the structure of user activity. Activity-based timelines might facilitate better interaction by providing a multiple scale view of the past. For example, selected thumbnails could be integrated with a timeline to represent individual episodes. Users could “zoom” in and out of an episode as well as more efficiently travel through time. These thumbnails could be selected and displayed based on landmarks in user activity as opposed to calendar time. Activity landmarks might be simple changes in activity patterns, such as changes in window arrangement, or applications, or a change in physical location. Using this type of low-level landmarks might be effective because their visualization promotes inference on the part of the user without the need for machine-based understanding of the activities. A selected segment based on a landmark could be played back to assist in reestablishing its context as well as aiding users’ memories for the specific meaningful details of the activity, what was being done and what was being used to do it. Combining visual timelines with access to views of episodic activity could be an efficient mechanism to support browsing large volumes of activity history data because of their ability to support multiple timescale and multiple memory cues.

Timeline representations that reflect the structure of work rather than just simple linear time seem particularly promising for supporting navigation as opposed to searching. But searching is useful when we actually recall the details of a specific episode from the past. Searching the past could be implemented in a toolbar fashion like Google Desktop or Spotlight. But there are important differences with traditional search to consider: how do we use activity-based parameters instead of text query terms? For example, one might initiate a query by specifying that she is looking for an episode that occurred on a Friday afternoon because she recalls she was working at home during the episode she wants to retrieve and she commonly works at home on Fridays. She might also remember that she had an email exchange with a colleague about this topic and that could be used to further refine the query. After getting back a number of episodes, she might use a thumbnail augmented timeline to select short segments to play. Upon finding the specific segment she might play it and be reminded of other documents and spreadsheets that were involved and that she will need when resuming the activity.

CONCLUSIONS

Researchers from many disciplines are taking advantage of increasingly inexpensive digital video and storage facilities to assemble extensive data collections of human activity captured in real-world settings. Hollan and Hutchins [17] have argued that the ability to record and share such data has created a critical moment in the practice and scope of behavioral research.

Examining the fabric of everyday work activity reveals a *dialectic* process between what people want to accomplish and how it can be accomplished within the constraints of the facilities and applications available. No single application can provide full support for getting back into past work, because most real-world activity involves the simultaneous use of multiple tools. As new technology develops, people adapt the tools to their tasks, and new work practices emerge, like using IM for inter-office communication. Although work activity from an application perspective can appear fragmented and disjoint, it is important to understand that workers perceived themselves as working on a *single* activity. More importantly, this single activity constitutes a temporal *episode* in memory that includes the interaction history of all the tools, people, and resources involved. Activity histories have tremendous potential to assist where single applications left off. Recording activity at the operating system level brings together the separate pieces necessary to reinstate full mental context of an activity.

REFERENCES

1. Hill, W. C., Hollan, J. D., Wroblewski, D., and McCandless, T. 1992. Edit wear and read wear. *In Proc. of CHI 1992.*
2. Rekimoto, J. (1999). Time-machine computing. *12th annual symposium on user interface software and technology.* Asheville.
3. Freeman, E., & Gelernter, D. (1996). Lifestreams: a storage model for personal data. *ACM SIGMOD Bulletin* , 25(1):80–86.
4. Gemmell, J., Williams, L., Wood, K., Lueder, R., and Bell, G. (2004). Passive capture and ensuing issues for a personal lifetime store. *Proc. CARPE 2004*. 48-55.
5. Krishnan A. & Jones S. (2005). TimeSpace: activity-based temporal visualization of personal information spaces. *Personal Ubiquitous Computing*. Vol. 9:46-65. Springer-Verlag, London.
6. Jones, W. & Teevan, J. (Eds.) *Personal Information Management*, 2007.
7. Sellen, A. J., Fogg, A., Aitken, M., Hodges, S., Rother, C., and Wood, K. (2007). Do life-logging technologies support memory for the past?: an experimental study using sensecam. *Proc. CHI 2007*. 81-90.
8. Czerwinski M., Horvitz E.(2002). An Investigation of Memory for Daily Computing Events. *Proc. British HCI Group Annual Conference.*
9. Bradburn, N. M. (2000). Temporal representation and event dating. In A. A. Stone, J. S. Turkkan, C. A. Bachrach, J. B. Jobe, H. S. Kurtzman & et al (Eds.), *The science of self-report: Implications for research and practice.*, 49-61. Lawrence Erlbaum Associates Publishers.
10. Conway, M. A. (1996). Autobiographical knowledge and autobiographical memories. In D. C. Rubin (Ed.), *Remembering our past: Studies in autobiographical memory* 67-93. Cambridge, UK: Cambridge University Press.
11. Brewer, W. F. (1988). Memory for randomly sampled autobiographical events. In U. Neisser, & E. Winograd (Eds.), *Emory cognition project conference*, 21-90. New York, NY, US: Cambridge University Press.
12. Huttenlocher, J., & Prohaska, V. (1997). Reconstructing the times of past events. In N. Stein. P. A. Ornstein, B. Tversky, & C. Brainerd (Eds.), *Memory for everyday events*. Mahwah, NJ: Lawrence Erlbaum Associates.
13. Huttenlocher, J., Hedges, L., & Prohaska, V. (1988). Hierarchical organization in ordered domains: Estimating the dates of events. *Psychological Review*, 95(4), 471-484.
14. Shum, M. S. (1998). The role of temporal landmarks in autobiographical memory processes. *Psychological Bulletin,* 124(3), 423-423.
15. Huttenlocher, J., Hedges, L., & Prohaska, V. (1988). Hierarchical organization in ordered domains: Estimating the dates of events. *Psychological Review*, 95(4), 471-484.
16. Ringel M., Cutrell, E., Dumais, S. & Horvitz E. (2003). Milestones in Time: The Value of Landmarks in Retrieving Information from Personal Stores. In: *Proceedings of IFIP INTERACT03: Human-Computer Interaction 2003*, Zurich, Switzerland. p. 184.
17. Horvitz E., Dumais, S. & Koch P. (2004). Learning Predictive Models of Memory Landmarks, *CogSci 2004: 26th Annual Meeting of the Cognitive Science Society*, Chicago, August 2004.
18. Chau D. H., Myers B. & Faulring A. What to Do When Search Fails: Finding Information by Association. *In Proc. CHI 2008.*
19. Hollan, J. D. and Hutchins, E. L. Opportunities and Challenges for Augmented Environments: A Distributed Cognition Perspective. In *Designing User Friendly Augmented Work Environments: From Meeting Rooms to Digital Collaborative Spaces*, Saadi Lahlou (Ed,), Springer, 2009, In Press.
20. Williams, M. Hollan, J. (1981). The process of retrieval from very long-term memory, *Cognitive Science*, 87-119.

rhythIMs: a visualization of patterns of online and physical presence

Judy Chen
Donald Bren School of Information
and Computer Sciences
University of California, Irvine
Irvine, CA 92697-3440, USA
judychen@ics.uci.edu

CHI 2009, April 4 – April 9, 2009, Boston, MA, USA
ACM 978-1-60558-247-4/08/04.

Introduction

Time and temporality play a large part in structuring our everyday lives, especially in our interactions with others. When people arrange social gatherings, all involved parties must agree upon a suitable time. Even impromptu gatherings can necessitate temporal sequences, e.g. "I'll meet you in the lounge in five minutes." As a long-term research goal, we are interested in exploring how awareness of other people's temporal patterns affects the ways in which people collaborate, coordinate and communicate. We are investigating the effect of this awareness on communication effectiveness and cohesion through the deployment of our system, *rhythIMs*, in which we present visualizations of instant messaging (IM) and physical presence on a semi-public situated display.

The potential for using a buddy's instant messaging presence as an indicator of social presence remains largely unexplored. Recently, studies have noted that a buddy's online presence can serve as a moderately accurate indicator of his social presence and availability. Several prototypes for online presence displays have been designed to focus on awareness and connectedness in [2]. However, these displays have largely been designed for individual use and provide up-to-date IM status changes for a single person.

Rather than focusing on presenting an individual user with IM presence information in real time, *rhythIMs* instead focuses on presenting a small group of people with visualizations that represent the group's IM and physical presence as it has built up over time. We are exploring the interactions that take place around such a display, both between the display and its users and amongst the people in the space itself, and are examining their implications for designing effective visualizations.

Related Work

A number of research projects have studied the temporal structures that exist in people's everyday work lives. Begole et al. analyzed computer activity and visualized the data over varying time-series. They found that patterns differed not only between individuals but also within individuals according to time of day, day of week and an individual's location [1]. Both Soylent [3] and Themail [8] present the user with visualizations of his email activity, revealing the social and temporal structures that exist in our communication with others and indicating the evolution of social relationships over time. The implications for this work are useful for revealing a shared sense of time amongst a group of people and for coordinating social contact for both co-located and distributed groups.

One method of communication that has become enormously popular over the last decade is instant messaging. As of 2004, it was estimated that 53 million American adults use instant messaging [7]. The rapid rise of instant messaging in both the workplace and recreational environments has sparked a number of studies on IM, many of which have focused on its role in supporting lightweight, informal communication. More recently, researchers have begun to note the advantages of using text-based point-to-point communication systems such as instant messaging in supporting awareness in a distributed group setting. In their study of instant messaging use in the workplace, Nardi et al. document that instant messages are frequently used for asking quick questions and clarifications about ongoing work tasks, coordinating and scheduling both formal and impromptu meetings, and keeping in touch with friends and family [6]. The use of instant messaging as a means of socializing, coordinating social congregations and collaboration has been reported in recreational contexts as well [4]. These studies show that IM supports awareness by providing users with a lightweight channel of communication that allows for the immediacy of face-to-face and over-the-phone interaction but without the overhead of maintaining these types of interactions.

rhythIMs

Our primary design objective is to support awareness and to achieve a feeling of connectedness for a group of people. By connectedness, we are referring to "a positive emotional appraisal which is characterized by a feeling of staying in touch within ongoing social relationships" [5]. We will do so by publicizing opportunities for informal and spontaneous communication and reflecting lightweight information about a group's instant messaging activity and physical presence back to the group. We reiterate that the visibility of one's instant messaging "buddies" is a lightweight means of supporting awareness in that the knowledge of a buddy's IM status can facilitate the spontaneous interactions that instant messaging enables and contributes to a feeling of connectedness.

While previous work has already been done to support this hypothesis, these systems are largely desktop-based and require a person to explicitly focus on an application to pull the awareness information it provides. Instead, we argue for pushing information out to people unobtrusively to minimize the effort required of them to access it. By moving awareness information off the desktop, we are also creating new opportunities for social and physical interaction.

In our visualization, each person is represented by a set of layers colored at three different saturation levels to differentiate the data for IM presence (33% saturation), physical presence (66% saturation) and for when a user is on IM while being physically present near the display (100% saturation). A group of people is then represented by the aggregation of all the layers stacked on top of each other in an additive fashion. The visualization is then split by a vertical line, which indicates the current time of day. The layers to the left of the line indicate the presence activity that has occurred so far within the past day. The layers to the right of the line indicate the patterns of presence activity that have aggregated throughout one's history. The result is a visualization that shows people what their recent IM and physical presence activity looks like (on the left) and predicts what their future activity for the rest of the day will look like based on their past history (on the right). The maximum layer height is a function of the total number of users so that each layer is allocated equal amounts of maximum vertical space. The actual thickness of each layer is calculated as follows: for each minute of the day, the height of a layer is the result of multiplying the maximum layer height by the ratio of the number of days the user has been logged on to an instant messaging service, or in the case of physical presence, the number of days the user has been detected near the display at that specific minute, to the number of days that the system has collected data for that individual. An example of what the visualization looks like for a single user is illustrated in Figure 1. We consider a user to be logged on when he is signed in to an IM service and offline otherwise. We are also using Bluetooth scanner and using the detection of a user's Bluetooth device such as his mobile phone or laptop as a substitute for physical presence.

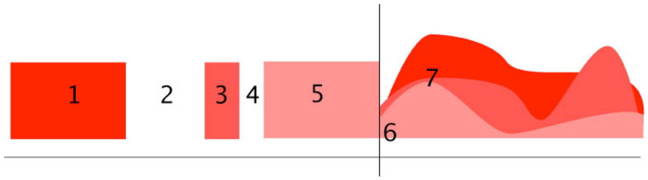

Figure 1: An example of the *rhythIMs* visualization for a single user. The regions are labeled as follows. 1: The user was online while physically present near the display. 2: They leave the room and sign off IM. 3: They re-enter the space but do not go online. 4: They leave the space again. 5: They go online but are working remotely, away from the display. 6: From their past history, they are likely to be online but working remotely over the next few minutes. 7: From their past history, they are as equally likely to be on IM near the display as they are to being on IM remotely.

The goal of the visualization is to give people an overall sense of their own IM and physical presence, as well as the temporal patterns of others around them. The idea of using layers to compare one entity with respect to a larger whole was inspired by the Artifacts of the Presence Era installation [9]. We added the height of each additional layer to that of the previous one to form an accumulation of layers. Like the gradual accumulation of layers in sedimentary rock, the cumulative view represents a historical buildup of information. This approach allows us to examine an individual's contribution with respect to the group. It also allows us to view the data at a higher level, where we can examine the layers as a whole to see what temporal patterns emerge from the collective as well as from the individual. The visualization builds on this notion of examining properties of an individual's behavior and how they compare and contribute to an emerging collective behavior. By layering each of the individual's contributions in an additive fashion, we are introduced to two dimensions of information: the patterns exhibited by the individual and how they contribute to the patterns exhibited by the collective.

Conclusion

We are currently in the early stages of deploying and evaluating *rhythIMs*. In our evaluation, we are examining the ways in which IM and Bluetooth presence can be a substitute for physical presence. Do people feel more connected to others through the visualization? How accurate are people's expectations of when others are physically present or on IM and what cues do they use? What factors lead to inaccurate expectations and what can be done to alleviate those inaccuracies? Our focus is not to empirically investigate *rhythIMs* as a tool for increasing awareness but rather, to use *rhythIMs* as an exploratory tool to better understand the process of visualizing temporal data, to examine what people do with that information and to explore the kinds of interactions that take place around such visualizations.

References

[1] Begole, J., Tang, J.C., Smith, R.B. and Yankelovich, N. Work Rhythms: Analyzing Visualizations of Awareness Histories of Distributed Groups. *Proceedings of CSCW'02*, ACM Press (2002), 334-343.

[2] Dey, A.K. and De Guzman, E.S. From Awareness to Connectedness: The Design and Deployment of Presence Displays. *Proceedings of CHI 2006*, ACM Press (2006), 899-908.

[3] Fisher, D. and Dourish, P. Social and Temporal Structures in Everyday Collaboration. *Proceedings of CHI 2004*, ACM Press (2004), 551-558.

[4] Grinter, R.E. and Palen, L. Instant messaging in teen life. *Proceedings of CSCW'02*, ACM Press (2002), 21-30.

[5] Ijsselsteijn, W.A., van Baren, J., and van Lanen, F. Staying in Touch: Social Presence and Connectedness through Synchronous and Asynchronous Communication Media. *Proceedings of HCI International*, PUBLISHER??? (2003), 924-928.

[6] Nardi, B.A., Whittaker, S. and Bradner, E. Interaction and Outeraction: Instant Messaging in Action. *Proceedings of CSCW'00*, ACM Press (2000), 79-88.

[7] Shiu, E. and Lenhart, A. How Americans Use Instant Messaging. *Pew Internet and American Life Project* (2004), 1-19.

[8] Viègas, F.B., Golder, S. and Donath, J. Visualizing Email Content: Portraying Relationships from Conversational Histories. In *Proceedings of CHI 2006*, ACM Press (2006), 979-988.

[9] Viègas, F.B., Perry, E., Howe, E. and Donath, J. Artifacts of the Presence Era: Using Information Visualization to Create an Evocative Souvenir. In *IEEE Symposium on Information Visualization 2004*, IEEE (2004), 105-111.

Product Usage Data Challenges: Volume, Detail, Abstraction

Lynn Cherny
Ghostweather Research & Design
209 Edgewater Drive
Framingham, MA 01702 USA
Lynn@ghostweather.com

CHI 2009, April 4 – 9, 2009, Boston, MA, USA
ACM 978-1-60558-246-7/09/04.

Abstract
Many software companies consider themselves "data-driven" in their operations, yet the data they collect is often under-leveraged. At one company, "Design Company," the temporal analysis component is generally missing, due to poor infrastructure for such analysis and lack of internal analytic expertise. Temporal data could be of enormous value in this and similar companies, given the right approaches to database design, algorithms, and front end visualization tools.

Keywords
Temporal, data abstraction, measures, analysis, data mining, graphs, databases, clustering, patterns, usability

Data In "Design Company"
Usually far from the research ideal, data collection inside business settings is nevertheless critical to operations. More and more companies consider themselves "data-driven" and chart their successes and growth with different metrics that include trend data presented in various ways. The popularity of the Net Promoter Score as a marketing measure of customer satisfaction [3] is just one metric currently in use in many companies, usually tracked carefully over time.

Within software development organizations, reports of bug counts of various severities constitute the main metric used to determine whether a product is ready to ship. Other metrics are increasingly being used to inform design processes and product quality measurements, such as mean time between crash reports, satisfaction and usability of specific features, performance measures (such as time for a system to complete operations), and product usage statistics [1, 2]. However, these data often live inside complex, historically ad hoc databases with poor or non-existent user interfaces; and such data are often not considered as stats that might be tracked and visualized *over time*. Reasons for this vary from inadequate infrastructure for performing temporal comparisons, impoverished understanding of how to interpret this data as a tracking mechanism of value to operations, and frequently a lack of resources for long-term corporate infrastructure projects.

My interest lies in learning ways to compute and present such data to provide insight for corporate decision processes, particularly where improvements in product quality and design are involved. Product usage data is of particular focus to me in this endeavor, in that it is often least utilized and is considered the most costly to collect, store, and analyze at “Design Company.”

Product Usage Data

In most engineering organizations, quantitative data is generally respected more than qualitative data, where “respected” means “more likely to be heard and to influence management behavior.” One type of quantitative information that has great potential for impact, but is not used to full potential, is data on product feature usage.

“Design Company,” like increasing numbers of software companies (and most web companies), instruments its own products for collection of usage data. This data is collected in an automated “opt-in” procedure, advertised as a method to “improve the product” by providing data for the development team. The aspect of this data that is currently most useful is crash reporting; analysis of crash frequency is a critical measure of quality in many companies [2] and Design Company is no different. Yet such reports at Design Company are produced via an antiquated VB script that produces multi-page Excel files distributed by email once per month. Modification to present richer interaction and long-term trend data is extremely difficult and requires dedicated developer attention to the old scripts.

Challenges of Micro Detail

Collected along with crash data are data on interface-level feature usage. This data lives at a low level of granularity, sometimes at the level of mouse movements. This micro detail frustrates most internal data consumers. Even an individual session contains too much detail to be useful in a general manner with the current internal tools; yet the promise of quantitative data would seem to be generalization across many sessions. Questions of great interest to the organization include: *In what contexts are various “help” commands used? Does usage change after introduction of new features? What features are NOT being used? Does crashing occur in certain contexts of use? What contexts cause the most repetitive “clicking”?*

Micro-level usage data pose several problems (see [1] for more):

- Volume of the data: It takes up a lot of room to store, especially at very low levels of granularity. The database administrator at "Design Company" is regularly asking to eliminate or reduce the data collected due to space issues and its lack of usefulness so far.

- Costs of analysis for a data-mining approach: It's computationally time-consuming and conceptually difficult to find the right algorithms to generalize from this data, particularly given the indirection in the database (commands are "coded" and multiple tables are required to identify them in text strings, and even then, the operations are not mapped one-to-one to the interface in a transparent way).

- Even for "simple" questions, the role of context, or temporal order, becomes critical: E.g., the length of time of a specific operation is easy to evaluate, but may only be valuable when looked at in context of surrounding operations. Adding this dimension becomes expensive, for the above reasons.

Sample Application One: Mouse Movements
Web site analytics packages present page click data in terms of button/link click heat maps. In a design application, the actual x,y coordinates of movement and clicking may be important, however. Here is an example of a flattened mouse movement patterns in the document window ofrom one user session, with "older" movements faded out relative to more recent movements:

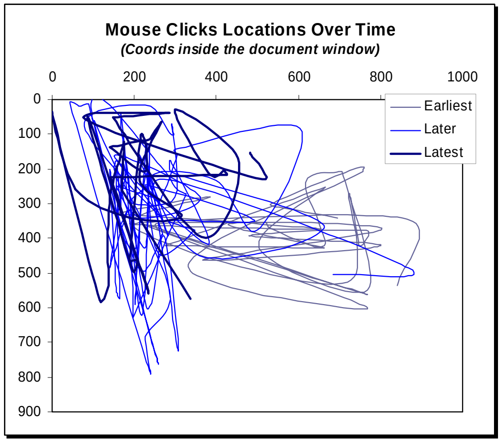

Figure 1: Mouse Clicking in Document Window (Traces)

Such diagrams suggest that the primary action in this product stays relatively close to the tool bar on the top left of the application window, although there are patterns of location by time that need further context analysis. The designers might consider moving buttons from the upper left to floating contextual toolbars, to reduce mouse travel to the left corner.

Sample Application Two: Performance Time
In Figure 2, a sequence of micro-operations shows long-duration operations as higher peaks. The same operations take different amounts of time, depending on when they were performed (in this case, probably due to increasing amounts of data in the document and decreased amounts of available memory):

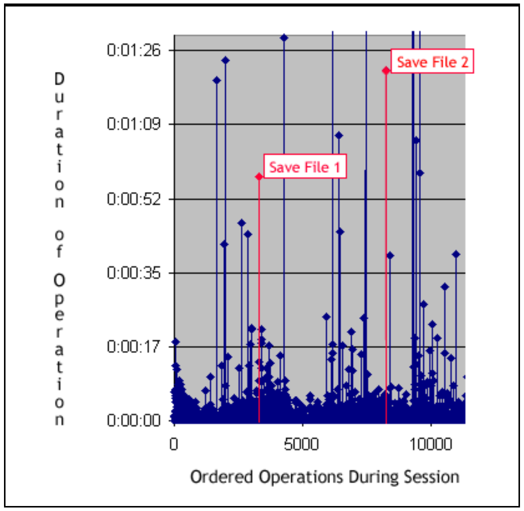

Figure 2: Durations for the Same Operation, at Two Different Points in a Session

A chart of memory usage over time in this session confirms that there is less available memory later in the session, when the user saves the second time.

Sample Application Three: Abstraction Needed
Different stakeholders require different levels of analysis of sequential operations. Clustering methods or pattern matching techniques might help analysts identify useful "meta levels" as suggested from fictional data representation in Figure 3 (also see [1]). Low-level operations may be grouped in different ways, and hierarchical patterns may emerge based on co-occurrence of command sequences or sets. As in other data mining problems, the analyst's role in interpreting the data from a position of domain knowledge is critical to drawing useful hypotheses and conclusions from the data. Tools for interactive analysis of this type are still lacking in "Design Company" and others with this type of detailed feature usage data.

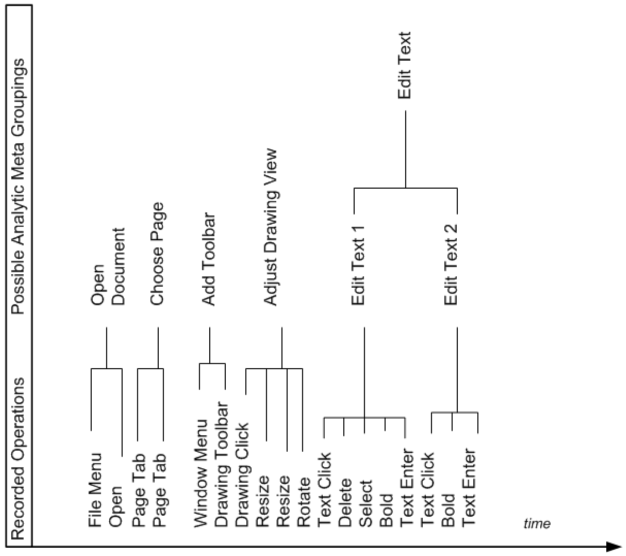

Figure 3: Analytic Groupings Potentially of More Use than the Micro-Recorded Operations

Citations

[1] Hilbert, D., and Redmiles, D. Extracting usability information from user interface events. *ACM Computing Surveys 32*, 4 (2000), 384-421.

[2] Jalote, P., Murphy, B., Garzia, M., and Errez, B. Measuring the Reliability of Products. From *Int. Symp. on Software Reliability Engineering (ISSRE04)*, Industrial Practices Track (Nov 2004).

[3] Reichheld, F. *The Ultimate Question: Driving Good Profits and True Growth.* Harvard Business School Press, Harvard, MA, USA, 2006.

Doing architecture on telecom data topography

Eunju Han
PhD candidate
Innovation Design Engineering
Royal College of Art
Kensington Gore
London, UK SW7 2EU
eunjuhan.arch@gmail.com
eunju.han@rca.ac.uk

CHI 2009, April 4 – April , 2009, Boston, the US

This design-based research aims to propose location-based urban metaverse related to socio-technological factors through investigating the relations between physical structure of urban space and temporal telecommunication behaviors.

Keywords
Locative metaverse, Telecommunication behavior, Ambient information, Data topography, Digital occupancy, Temporal communication data

ACM Classification Keywords
Future of Virtual Reality, Human Interaction, Teleimmersion

Short Proposition Paper
Everyday mobile phone use is highly connected with urban life, which is generating data that include spatial and temporal information reflecting people's urban daily life. In the meantime, as the groundings for spatial data, our built environment has consisted of layers reflecting time of sociology and history in a society. Also, urban telecommunication is producing temporal data along urban spatial axis for occupancies and movements. Through the temporal data juxtaposed with deconstructed urban layers, we may see teritorialization of telecommunication in urban space, which is used as a spatial structure of urban metaverse

that is virtual and locative urban layer(s) where we can utilize communicative places.

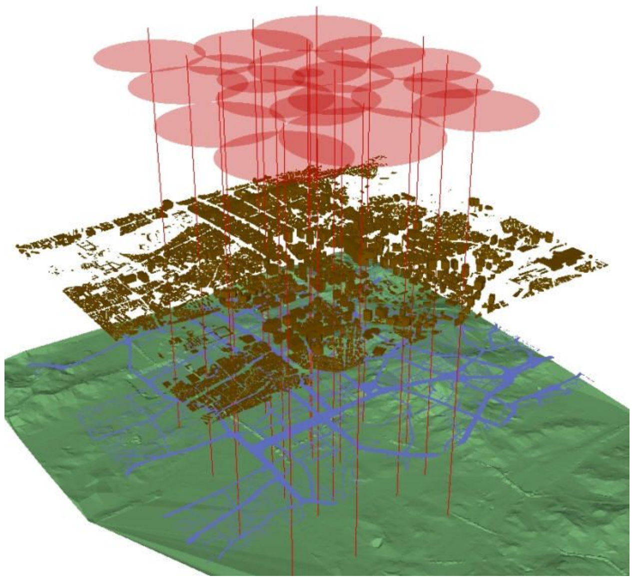

Figure1. Telecommunication cell spaces juxtaposing deconstructed urban layers; natural topography, roads and streets, solids and voids by buildings (from the bottom). The ends of red lines indicate the center of cell spaces and locations of base stations.

It will be investigated with the idea of this geological concept to make temporal information compatible with spatial data; 4-dimensional data topography in real urban map with the numbers of mobile phone calls dependent on time variations. To be specific, the spatial structure of urban telecommunication is revealed by a hybrid mapping between urban physical layers such as natural topography, heights of buildings, and networks of roads and 3-dimensional data contour; deconstruting, layering and juxtaposing. And then, applying temporal concept conveys 4-dimensional data moving topography that is the rhythm of telecommunication behavior and its territorialization in urban space. This is the ground to identify a new cognitive communication domain of urban space and, at the same time, provide a concept for designing locative metaverse, reflecting on telecommunication behavior.

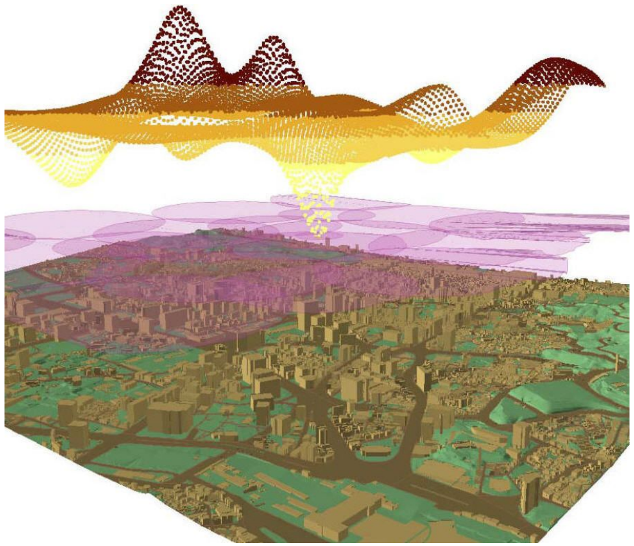

Figure2. Telecommunication cell spaces juxtaposing deconstructed urban layers; natural topography, roads and streets, solids and voids by buildings (from the bottom). The ends of red lines indicate the center of cell spaces and locations of base stations.

Eventually, locative metaverse will be generated at the level of built environment and at the level of interface of telecommunication devices. At urban space, architectures may be responsive depending on mobile scribers' behaviors and arbitrary and/or intended controls by telecom devices. At interfaces reflecting real spaces in mobile phone, we may virtually create customized communicative spots and then do locative communication with others in real time through putting annotations and browsing messages made of images, sounds, touches as well as texts and voices.

As contributions, urban media based on location may deliver a new urban communication through making us utilizing the traces of people's urban activity by depicting as informative pictures. At private individuals' level, it allows to use built environment as another media for communication. At local authorities' level, it may be made use for public awareness systems and transportation locally and even widely. At companies' level, retail strategies and ads for focus groups connecting with local areas will be available in real-time through locative media. Finally, the "ubiquitous" city will work.

Challenges

For interaction with temporal data in this research, some works for compatibility are needed. Accordingly, data visualization along urban real map was accomplished through GIS works with the idea of a topographical concept.

In my opinion, temporal data, especially regarding telecommunication in urban space, reveal a rhythm of spatial occupancy (and/or use). Also, the (concept of) rhythm may tell us attributes of spaces and places as media. In this respect, I would love to discuss about how to decode temporal data as a genetic factor pertaining to space, place, and intangible domain.

MemTable – Embodying Contextual History in Shared Workspaces

Seth Hunter
Workshop Applicant
MIT Media Lab
Fluid Interfaces Group
20 Ames St, E15-313
Cambridge, MA 02139 USA
hunters@mit.edu
http://fluid.media.mit.edu

Pattie Maes
Advising Project
MIT Media Lab
Fluid Interfaces Group
20 Ames St, E15-313
Cambridge, MA 02139 USA
pattie@media.mit.edu

Abstract
This position paper describes current research questions and challenges we have encountered as we work to design a table with a memory. The goal of the system is to augment work and meetings on a shared tabletop surface by providing a contextual interface for recording, annotating, reviewing, and reflection by the users of the table. Existing tabletop interfaces introduce new interaction techniques but neglect serious inquiry into the integration of these features with the long-term work practices. Our goal is to introduce a tabletop for literal and metaphorical reflection, contextual retrieval of information, and documentation of process. MemTable examines how an interface that embodies the history of its use can be incorporated into our daily work lives.

Introduction
Finding an appropriate context for the long-term use of an interactive table requires designing a table with features that augment, support, and document how that community is using the surface in the temporal and physical domains. We are designing a table that supports three modes of engagement: (1)multi-modal input for documentation and recording, (2)personalized review of the history, and (3)visualization of long term usage.

At the Temporal Interaction Workshop on April 4, we hope to discuss the challenge of how to identify the most contextual and significant data during table use and customize that data for each person who returns to the table. What visualizations are most useful to help us reflect on our long-term work practices, and what form should those visualizations take? Currently we have begun prototyping and development of the memory table, some of which we present in this document, and we hope to have more developed work in progress to present and get feedback about during the workshop in April.

Project Description
The implementation of the MemTable consists of: the physical build and installation, and designing multi-touch software for reviewing and visualizing the temporal history. MemTable is designed as a large work table for use in the MIT Media Lab. We have tried to choose a context that could serve as a template for other small workgroups such as design studios, conference meeting rooms, and organizations.

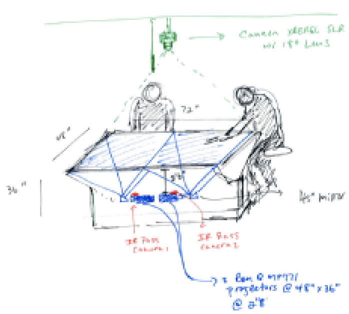

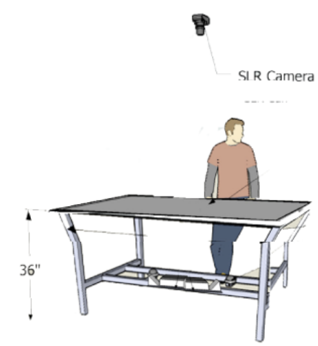

Figure 1. Preliminary drawings of the dimensions and technology in the proposed table

The table can record audio, high resolution images of the surface, and pen based annotations made by the users. It is also equipped with RFID readers so that users can tag recorded information by a unique ID.

We hope to support meeting scenarios like the ones pictured in figure 2, capturing temporal data about the objects on the surface, the audio in the space, and who is at the table.

Currently we are working on the technical aspects of developing the project, and mocking up scenarios where the temporal history become can be accessed "Just in time" while a meeting is happening.

Figure 2. Example Scenarios of Use with Objects

Reviewing Temporal Data

What are the best metaphors for reviewing temporal data? We have begun with a rotating spiral as a metaphor for unwinding from that past into the present – with layers within the spiral corresponding to different timestamps of types of data, drawings, ID, and audio.

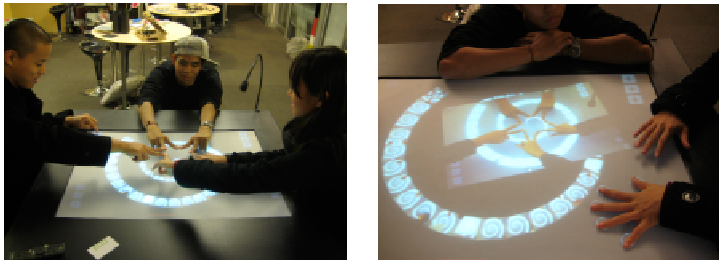

figure 3. A quick multi-touch prototype, one method of exploring a history of images we have explored.

The Temporal Interfaces workshop could provide an opportunity to present our current progress, and explore other possibilities for visualization of temporal events and long term use, such as the work by Frenanda Viegas in PostHistory, and Martin Wattenberg in History Flow pictured in figure 4.

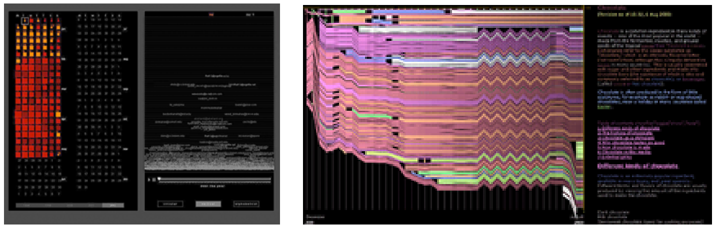

figure 4. A time based, and linear visualization of historical data by Viegas, Wattenberg.

In addition, we see an opportunity to visualize the social relationships between people in the space, based on their "time of use" patterns around the collaborative table such as in figure 5 below.

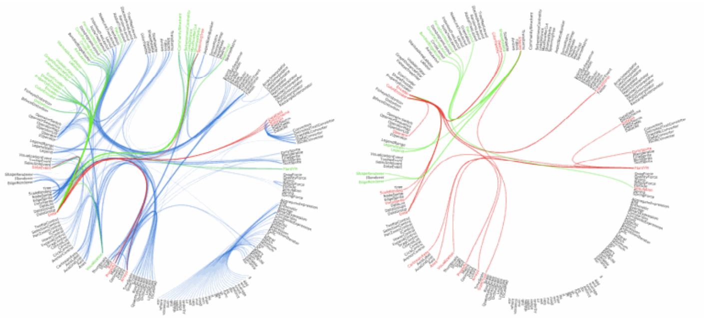

f
igure 5. An example of visualization techniques from the flare Libraries developed at UC Berkely. Here we look at all the people who have visited the table and their connections to each other.

Conclusion

In Doland Schon's influential book, "The Reflective Practicititioner"[2], he emphasizes the need to bring reflection to the center of an understanding of what we do, looking at our experiences and building new understanding from the data we gather.

Just as physical objects embody the temporal expressions of our physical presence (coffee stains, wear, notes) – digital objects have the potential of being integrated into our physical space and capturing the contextual data generated in the space. These digital objects can then embody that information by presenting it in the same special context as it was generated.

Great potential exists for other surfaces to record data about how they are used and adapt to augment the work practices of the people using them, but many questions need to be answered about how that data should be visualized, and what the most useful interfaces for navigating that information are. We hope you will consider us for participation in the workshop, we were very excited by the specific theme of temporal interaction and how well it correlated to the MemTable project.

For further information about this work please see: http://fluid.media.mit.edu or contact hunters@mit.edu

Citations

[1] Martin Wattenberg, Fernanda Viegas http://www.bewitched.com/historyflow.html

[2] Schön, D. The Reflective Practitioner: How Professionals Think in Action, Basic Books, New York, 1983

Temporal data visualizations for Air Traffic Controllers (ATC)

Christophe HURTER
DGAC DSNA DTI R&D
7, Avenue Edouard Belin
31055, Toulouse France
christophe.hurter@aviation-civile.gouv.fr

Stéphane CONVERSY
ENAC
7, Avenue Edouard Belin
31055, Toulouse France
stephane.conversy@enac.fr

Jean-Luc VINOT
DGAC DSNA DTI R&D
7, Avenue Edouard Belin
31055, Toulouse France
Vinot@cena.fr

CHI 2009, April 4 – 9, 2009, Boston, MA, USA
ACM 978-1-60558-246-7/09/04.

Introduction

Fundamental research in visualization is concerned with the impact of presentation on visual perception and understanding [5] [6].

In current Air Traffic Control (ATC) environments, Air Traffic Controllers (ATCo) use several visualization systems: radar views, timelines, electronic strips, meteorological views, supervisions etc… The information displayed is intrinsically temporal: aircraft positions that evolve over time, conflicts between aircraft at a given time, airspace congestions… Each of these visualizations is rich and dynamic: it displays numerous visual entities that move and evolve over time. Furthermore, this considerable volume of information must be understandable with a minimum cognitive workload. As traffic increases and safety criteria become more demanding ATC requires this new kind of visualizations.

Our research focuses on the production of new *efficient* visualizations with temporal data. We characterize a more *efficient* visualization as one in which a greater volume of information can be perceived and understood with a smaller or equivalent cognitive workload (while at the same time reducing the error rate in the perception of the information)

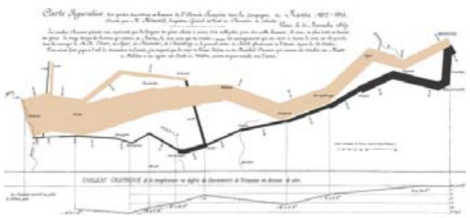

Minard's visual history of Napoleon's Russian Campaign

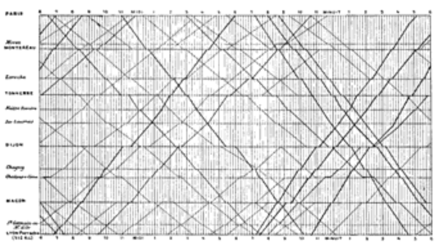

This diagram, showing the time schedules of trains from Paris to Lyon and vice-versa, was published by Marey in his book.

Time and visualizations

Bertin [1] was the first to define representation rules that produce a non ambiguous representation: the semiology of graphics. This semiology is a complete system composed of signs and rules which permit the building of planar maps or diagrams. This system is a powerful tool to represent, memorize, understand and communicate information. It describes and explains perceptual phenomena and properties underlying the act of reading an abstract graphic. Bertin also introduces the variables that are used to code information visually: position, size, orientation, color. Bertin did not provide special rules for time representation. For us, time is just another data attribute in a dataset. However, noteworthy representations of temporal information do exist (Train schedule string-line graphics by Marey, Minard's visual history of Napoleon's Russian Campaign…). These graphics address a major event to show temporal data with a static representation. These views are not a simple time series (or timeline). In past works, we conducted a precise analysis designed to understand the temporal dimensions more fully. We used Bertin's work and we tried to characterize design choices [2]. Then, we tried to compare design choices with objective criteria [3]. Finally we sought new visualization to improve image efficiency

In this paper, we propose to illustrate one visual design that codes temporal information: The radar comet used by ATCo.

RADAR Comet:

It is very difficult to create a new design without any support. Nevertheless, we identify several different methods when building visual entities:

- An empirical approach: design based on trial and error methodology,
- A historical approach: design based on the continuity of previous work with a concern for adaptation to the given context,
- An ecological approach: design based on the respect of both physical and perceptual human factors
- A technological approach: design based on technological opportunities.

The design we propose to investigate is the Radar Comet. It instances all the mentioned sources of design. The first design source is the ecological design. A comet's visual properties were used for the first time in the early seventeenth century by Edmond Halley [4] (Figure 2) who coded the trade wind direction on a map. He coded the flow with a stroke. The comet has accurate design properties; it displays the direction of the shape and its tendency. The comet is composed of a bigger part, its head, and a smaller, its tail. Its head indicates the comet heading. The tendency indicates the direction. The curvature of this shape indicates if it is turning right or left and the degree of veering. This design is ecological in so far as it stems from our basic understanding of motion (drops in a puddle, shooting stars…). An ecological design is easily understandable

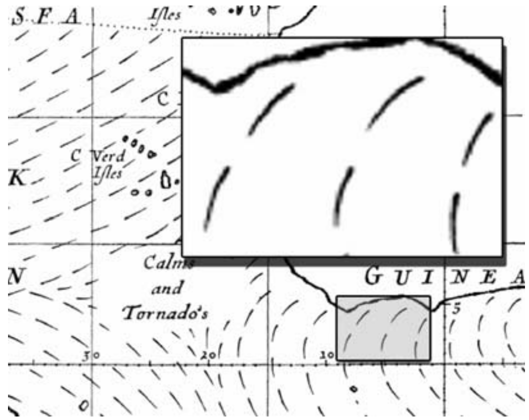

Figure 2: Halley's chart of the Trade Winds 1686.

ATC visualization derives some benefit of this comet. The design of the first radar comet is historical and technical. It is not directly based on the Halley design but on early radar equipment which relied on scope persistency (Figure 3). Old radar scopes retained the previous plot positions with the fading of the screen phosphor (ecological and technical design: the plot has a lifetime, the dot size and luminosity decrease over time). This kind of design has the same remarkable properties as the Halley-style comet: it displays the aircraft trajectory's curvature tendencies and shows if an aircraft is turning and the degree of veering. ODS is the main French radar view for the air traffic controllers. Its main goal is to display aircraft positions and to help ATCo to space aircrafts beyond the security minima. The radar track presents the aircraft positions, its speed (speed vector), name, altitude and speed as text (label). The design of the comet is built with squares, whose size varies with the proximity in time of the aircraft's position: the biggest square displays the latest position of the aircraft, whereas the smallest square displays the least recent aircraft position. Thus, in Figure 4, the shape of the comet indicates that the plane has turned 90° to the right and that it has accelerated.

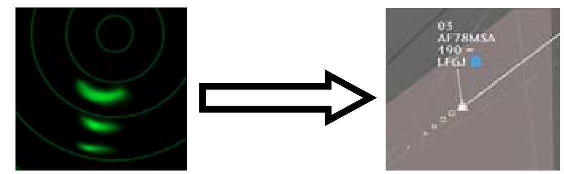

Figure 3: the shape decreases in intensity over time on a scope (left picture). ODS comet metaphor (right picture).

The positions of the aircraft merge through the effect of Gestalt continuity, in which a line emerges with its particular characteristics (curve, regularity of the texture formed by the points, etc). Furthermore, Gestalt continuity allows an occlusion resistance. This means that if there are many overlapping comets you can still distinguish each one.

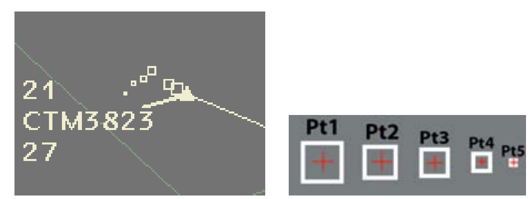

Figure 4: the French radar screen

The information coded by the comet is summarized in the following table. Italic script represents emerging information. For instance, the aircraft tendency data is not needed to draw the comet, but the comet nevertheless codes this information. The emerging process stems from the embedded time in the RADAR plot positions. The time can be easily derived into speed and acceleration.

The comet design is an efficient visualization of data time. The fact that the underlying data is temporal helps the reader to interpret naturally emergent information such as speed, tendency and acceleration.

ODS coded information	Visual code
Aircraft position	Position
ageing of each position	Size
Aircraft speed	Size (comet length)
Aircraft tendency (left, right)	Comet curvature
Aircraft acceleration	Regular/irregular point spacing
Aircraft entity	Gestalt (proximity and size)

Conclusion:

In this paper we depict only one simple visual design used by air traffic controllers but it remains noteworthy through its numerous roots and the impressively large amount of coded information. This design illustrates the emerging information process. ATC use several other interesting temporal designs.

The comet design addresses some issues: the design is efficient because it displays more information than the simple aircraft positions (emerging information). Furthermore, this design is easily understandable because is ecological (metaphor of the phosphor persistency, and the comet design).

We should be very interested in participating in any workshop focussing on temporal data, especially if it addresses representation issues. The perspective of a special journal issue is also a promising project. The representation of temporal data is a part of the work of one of our Ph.D students. Therefore, we should be eager to take part and to benefit from the opportunity to contribute our expertise and our visualization experience.

References

[1] J. Bertin , Graphics and Graphic Information Processing deGruyter Press, Berlin, 1977.

[2] C. Hurter, S. Conversy: Towards Characterizing Visualizations. DSV-IS 2008:

[3] C. Hurter, S. Conversy, V. Kapp, An Infovis approach to compare ATC comets ICRAT 2008

[4] J. Norman, W. Thrower, Edmond Halley as a Thematic Geo-Cartographer, Annals of the Association of American Geographers, Vol. 59, 1969.

[5] E.R. Tufte, The Visual Display of Quantitative Information, Graphics Press, Chesire, Connecticut, 1983.

[6] C. Ware, Information Visualization, perception for design, Morgan Kaufmann, 2002.

Viewing Personal Data over Time

Sonja Knoll, Aaron Hoff, Danyel Fisher, Susan Dumais, Ed Cutrell
Microsoft Research
1 Microsoft Way, Redmond, Washington. USA 98052
{sonjak, aaron.hoff, danyelf, sdumais, cutrell}@microsoft.com

ABSTRACT
Desktop search applications are changing the way we interact with personal information: we have a choice of whether to view files within their traditional siloed hierarchies, or brought together in search results. In this position paper, we discuss the advantages of temporal ordering of desktop search results, and present an interface that aggregates search results over time.

Author Keywords
Desktop search, personal information management

INTRODUCTION
Desktop search applications are changing the way we interact with personal information. Traditionally, personal computer systems kept information siloed and hierarchical: email lived only within email clients, and could only be read from within them; files lived within folders, and could only be accessed from desktop applications. While limited searches were available, they were inconvenient and difficult.

The advent of desktop indexes changed that. With a background process generating and maintaining a desktop index, and a series of indexing plugins that can extract words from all files, the file system is conceptually flattened: keywords can be used to collect documents from anywhere.. A user can search for content through email, word-processing documents, notes, spreadsheets, and other sources all at once. Indexes can also store metadata attributes; it is often possible to search for emails that contain attachments, sent before a particular date, in HTML format, or from a particular person—all information that is stored in metadata, rather than content.

This is far from news. The Stuff I've Seen project [3] and Phlat [2] have discussed the importance of exposing this index, and have presented tools that allow users to work with their personal information.

Most consumer applications—such as Google's Desktop Search and Microsoft's Windows Desktop Search—present a ranked list of hits or a rich list in response to search queries. The ordering is sometimes determined by relevance score and other times, it is sorted by a time attribute. Relevance rankings are traditionally timeless and based on text analysis; something I worked on several years ago is ranked in the same way as something from today. Even ordering by time is complex: as we have previously noted [3], there are many notions of "time" on a computer system and for users; files can be meaningfully labeled by the time they were created, last modified, or last accessed. We also know from research in cognitive psychology that people remember information not in terms of exact time, but rather in episodes and relative to other events in time [3]. It is critical to get all this right in order for the user to find the information they are looking for.

In this position paper, we discuss temporal aspects of desktop search results from two perspectives: mining novelty and information visualization. In each, we discuss several related projects and then present our own contribution.

Indexing and the Desktop
In order to access this information, the classic desktop search index needs to be enhanced. It is difficult, using current versions of desktop search, to directly ask questions like "has there been a recent burst of activity around a particular term?" Instead, every document that mentions the term would need to be retrieved, analyzed, and its metadata examined.

We have modified Windows Desktop Search to calculate a broader set of corpus statistics, to add new properties to the index, and to generate a forward index that can be used to later compute additional metadata. Both of the applications we discuss in this section are based on these modifications to Desktop Search.

PERSONAL VISUALIZATION AND TEMPORALITY
Showing richer temporal information can be useful, especially for personal content. A number of research projects have approached this question.

Research Approaches
LifeLines [10] is a visual representation of a user over time. Collected from specific sources, LifeLines can then show personal histories. LifeLines is designed for cases when rich metadata is available with extensive labeling that can divide events into categories and subcategories with beginnings and endings: health care records are used in their example.

LifeStreams [5] puts all documents in sequential order, allowing users to scroll through their documents through time. Integrated search and metadata selection are filters on the data.

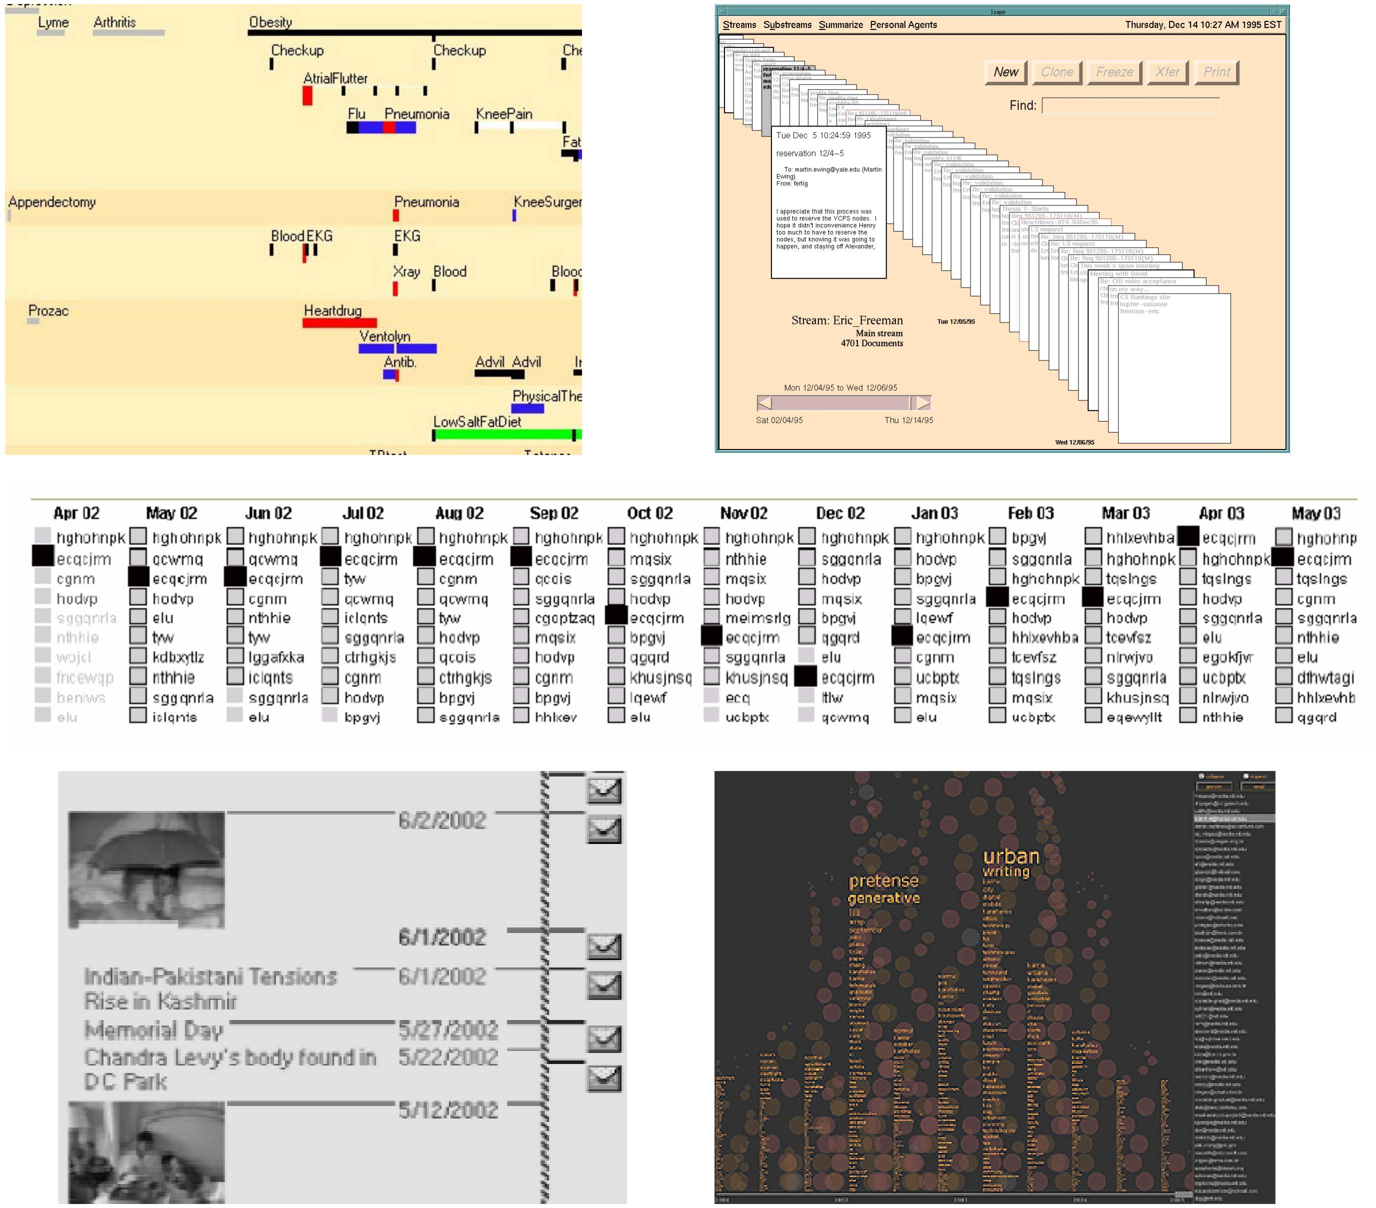

Figure 1: Five interfaces for personal information over time: LifeLines, LifeStreams, Soylent, Milestones, and Themail.

“Milestones” [11] added to the desktop search listing by annotating the search results with a timeline. Times on the result list are annotated by entries from the user’s own calendar (meetings, important events), photographs, etc, as well as collective information.

Limited to the email domain, Soylent [6] includes a temporal visualization that shows the top ten email correspondents in each month. The changes in that list over time can suggest different projects, priorities, and social circumstances.

Similarly, Themail [12] shows the most common words appearing in email over time. It allows users to see the general content of their email changing over time, and to look for major changes in their email at a glance.

Showing Personal Narratives

Personal Narratives is a visual tool that also shows information flows over time. A screenshot of Personal Narratives is shown in Figure 2. Like other temporal PIM visualization projects, PN looks at changes in personal information flow over time. Unlike these other approaches, however, it does not attempt provide an overall view of the dataset, but rather shows a single query term at a time.

Personal Narratives is a revision of *Narratives* [7]. *Narratives* visualizes change in the number of blog entries posted on news topics over time. In Personal Narratives we instead visualize the number of references to a single name or keyword over time.

Personal Narratives is driven by our modified version of Windows Desktop Search. We extract a random 200 words from each document, and create a time-labeled forward index based on them; we then push this forward index into a database and visualize them over time, showing the relative number of hits.

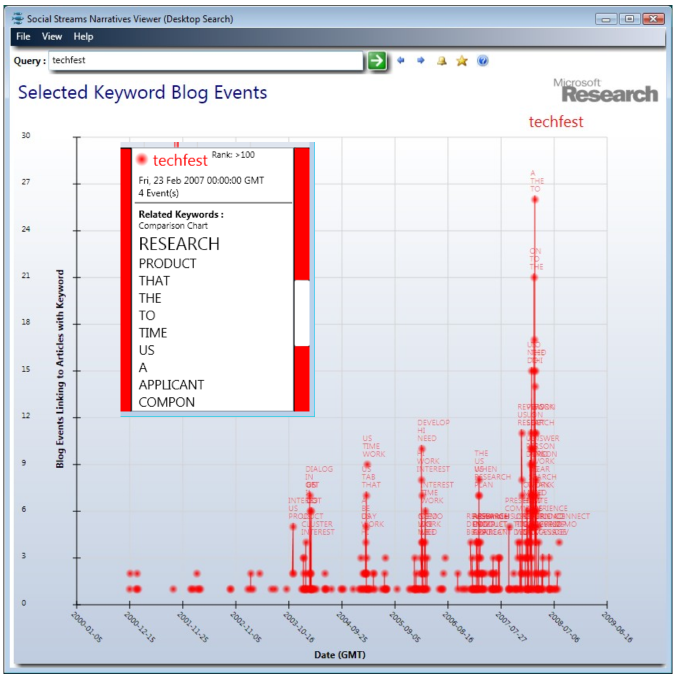

Figure 2: Personal Narratives, searching for the term "techfest". Inset, a tooltip from 2007.

In Figure 2, we show the interface for Personal Narratives. The screenshot covers the personal archive for one user over eight years: the timeline at the bottom runs from 2001 through 2008. Techfest is an annual event at Microsoft; thus, the red spikes occur once a year, building up in the email bursts coordinating techfest, then falling off after the event is over.

Narratives also finds words that closely co-occur with the source term. For example, if most documents that mention "client" also tend to mention "email", then a search for "client" would have "email" suggested. In Figure 2, we see that the more connected words in 2007 with "techfest" are "research" and "product". (Other terms in the list are less significant for this discussion.)

CHANGES IN PERSONAL INFORMATION OVER TIME

The visualization from Narratives inspired us to consider ways that information changes over time. The general research area of topic detection and tracking examines ways to handle text that has changed over time.

Kleinberg [9] has looked at temporal clustering, extracting meaningful temporal clusters of topics, including in email. Their work reports that personal email seems to cluster well into topics. Kleinberg extracts topics from ACM article titles (finding trends in terms like "web crawl" and "relational database") and from email.

Adar et al [1] looked at how users revisit web pages, and how websites change over time. They found that web pages change frequently and that user behavior around web pages was characteristic both of the users and of the pages. Some pages were revisited often, while users revisited other pages very infrequently.. By analogy, we might expect some PIM topics to have characteristic frequent, regular, or irregular patterns.

NewsJunkie [8] attempts to provide a personalized view of incoming news stories. Users can select a series of topics they are interested in, and it looks for novel stories on that general topic.

Personal Search Gadget

The Personal Search Gadget is similar to NewsJunkie, in that it locates novel material over time. It monitors a user's activity, looking for words and concepts that are new relative to a background model, such as new information compared to the entire index, or words that are new compared to the previous week.

In Figure 3, we show a screenshot from the Personal Search Gadget. The application is packaged as a small desktop "gadget" for Windows Sidebar, and so is peripherally visible as the user continues with other tasks. It shows the top four words that are distinctively prominent in the user's email. In this particular case, the user had been reading about weed killers, and thus both "glyphosate" and "roundup" have come to the top.

The user can dive in deeper to view recent and historical hits for any terms that the system surfaces; a click on the interface summons a search box to find all hits for these terms.

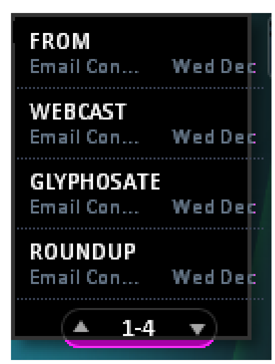

Figure 3: Desktop Search Gadget

CONCLUSIONS

A temporal perspective can be usefully applied to personal information; indeed, arguably, a temporal ordering is one of the most meaningful ways to approach personal information.

In this paper, we have reviewed several different directions that PIM research has followed for temporality, and we have discussed both a visualization and an adaptive desktop gadget that attempt to surface aspects of this temporal information. Personal Narratives attempts to visually explore topics over time, while the Personal Search Gadget looks for statistically unusual terms that have recently emerged. These, as well as other projects based around our enhanced Windows Desktop Search, will allow us to provide users with tools that allow them to better interact with their personal histories.

ACKNOWLEDGEMENTS

REFERENCES

[1] Eytan Adar, Jaime Teevan, and Susan Dumais. "Large Scale Analysis of Web Revisitation Patterns", CHI 2008, Florence, Italy, April 5-10, 2008

[2] E. Cutrell, D. C. Robbins, S. T. Dumais and R. Sarin (2006). Fast, flexible filtering with Phlat - Personal search and organization made easy. In *Proceedings of CHI'06.*

[3] Davies, G. and Thomson, D., eds. (1988). Memory in Context: Context in Memory. Wiley: Chichester, England.

[4] S. T. Dumais, E. Cutrell, E., J. J. Cadiz, G. Jancke, R. Sarin and D. C. Robbins (2003). Stuff I've Seen: A system for personal information retrieval and re-use. In *Proceedings of SIGIR 2003*.

[5] Fertig, S., Freeman, S., Gelernter, D. Lifestreams: An Alternative to the Desktop Metaphor. CHI 1996.

[6] Fisher, D., Dourish, P. Social and Temporal Structures in Everyday Collaboration. CHI 2004.

[7] Fisher, D., Hoff, A., Robertson, G., Hurst, M. Narratives: A Visualization to Track Narrative Events as They Develop. VAST 2008.

[8] E. Gabrilovich, S. Dumais, and E. Horvitz. Newsjunkie: Providing personalized newsfeeds via analysis of information novelty. In www13, 2004.

[9] J. Kleinberg. Bursty and Hierarchical Structure in Streams. Proc. 8th ACM SIGKDD Intl. Conf. on Knowledge Discovery and Data Mining, 2002.

[10] Plaisant, C., Milash, B., Rose, A., Widoff, S., Shneiderman, B. LifeLines: Visualizing Personal Histories. CHI 96

[11] Ringel (Morris), M., Cutrell, E., Dumais, S., and Horvitz, E. Milestones in Time: The Value of Landmarks in Retrieving Information from Personal Stores. Proceedings of Interact 2003, pp. 184-191.

[12] Viegas, F., Golder, S., Donath, J. Visualizing Email Content: Portraying Relationships from Conversational Histories. CHI 2006.

Visualizing Finger/Pen-Gesture Recognition in a Space Time Cube

Per Ola Kristensson
Cavendish Laboratory
University of Cambridge
JJ Thomson Avenue
Cambridge CB3 0HE
United Kingdom
pok21@cam.ac.uk

CHI 2009, April 4 – 9, 2009, Boston, MA, USA
ACM 978-1-60558-246-7/09/04.

Abstract
In this paper I suggest space time cube representation as a suitable representation for visualizing finger/pen-gestures. I then present a straight-forward geometric gesture recognition algorithm, and show how the behavior of this particular algorithm can be visualized to users in space time cube representation.

Keywords
Information visualization, visualization, spatiotemporal data, gesture recognition, algorithm visualization

ACM Classification Keywords
I5.1. Pattern recognition: Implementation – *interactive systems*.

Introduction
The space time cube is a 3D visualization technique for spatiotemporal data (originally proposed by Hägerstrand [1]). In space time cube representation, 2D spatiotemporal data is revealed to users by mapping the time dimension to one of the axes in a 3D cube. Space time cube representation tackles the fundamental problem of how to effectively convey complex spatiotemporal patterns to users. Unlike traditional 2D representations (for 2D spatiotemporal

data), space time cube representation conveys the temporal aspects of the data at a glance [2].

I argue that finger/pen-gestures are naturally represented in a space time cube. These gestures are articulated by the user on a 2D touch-sensitive surface and transformed into a sequence of time stamped sample points. The components in each sample point: the *x*, *y* and timestamp data, naturally maps directly into a 3D coordinate system.

As an example of space time cube visualization for finger/pen-gestures consider the following example. The user has pushed down a finger onto a touch-sensitive surface and is sliding the finger in a repeated circular motion, such that the motion eventually forms overlapping circles. Thereafter, the user lifts up the finger.

Using 2D representation this finger trace will be revealed to the user as a set of overlapping circles.

Now consider visualizing the same motion in a 3D space where the *x*-*y* axis is the Euclidean input space and the *z*-axis represents time (choices of axes are arbitrary). In this space time cube representation, the user's gesture forms a helix that immediately reveal to the user the dynamic nature of the gesture. Furthermore, if the user slowed down at particular segments in the original motion, this temporal property of the data is directly reflected in the space time cube representation, since the torsion of the helix changes as a function of the velocity of the original circular motion. This example demonstrates two possible advantages of space time cube representation:

1. It provides users with a bird's eye view of spatiotemporal data that enables users to quickly recognize global spatiotemporal patterns.

2. It enables users to detect complex spatiotemporal local variations at a glance.

It has long been hypothesized that space time cube representation is advantageous for spatiotemporal data analysis, particularly in revealing complex spatiotemporal patterns to users. A forthcoming paper [2] provides empirical evidence that this hypothesis indeed holds true: among other things, users do understand complex spatiotemporal patterns significantly faster using space time cube than using a baseline 2D representation.

The result in [2] paves the way for using space time cube representation for visualizing non-trivial spatiotemporal relationship. As an example, consider gesture recognition. An unsolved problem in gesture recognition is how to explain to gesture designers how conflicts in gesture set occurs, and how these conflicts relate to the gesture recognizer [5]. Below I illustrate a possible solution to this problem. This solution is based on visualizing a finger/pen-gesture recognition algorithm using space time cube representation.

Example: spatial finger/pen-gesture recognition and visualization

Below I explain how the finger/pen-gesture recognition algorithm works. Note that this recognizer is (crucially) very different in its nature in comparison to the one used in the gesture design toolkit in [5]. The algorithmic description below is to some extent adapted from Chapter 4 in [3] (first described in [4]), although in [3, 4] it is used in a different context than here.

Finger/pen-gesture recognition algorithm

Assume the user has articulated an input gesture $G = (\mathbf{g}_1, \mathbf{g}_2, \ldots, \mathbf{g}_k)$, with k ordered 2D sample points $\{\mathbf{g}_i\}_{i=1}^{k}$ in a Euclidean space.

Now, let $X = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n)$ represent an equidistant resampled version of G with n sample points, and let $Y = (\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_n)$ represent a template gesture (taken from a gesture set Ω), also equidistant resampled into n sample points. To recognize X we need to find the best matching template gesture Y^* in Ω.

When recognizing finger/pen-gestures it is desirable to be invariant of scale and translation displacements, since users cannot typically be expected to gesture in a predefined scale at a predefined position in the user interface. To achieve a normalized comparison between X and Y define an affine transform matrix $\mathbf{T}$ in homogeneous coordinates that transforms X into the same coordinate system as Y (corrected in, for example, scale, translation and rotation).

As a simple example, the matrix below scales and translates any gesture:

$$\mathbf{T} = \begin{bmatrix} s & 0 & dx \\ 0 & s & dy \\ 0 & 0 & 1 \end{bmatrix}, \tag{1}$$

where dx and dy is the distance between the two geometric centroids of X and Y, and s is a scaling factor determined as:

$$s = \frac{\max(w_x, h_x)}{\max(w_y, h_y)}. \tag{2}$$

where w_x, w_y, h_x and h_y refer to the widths and heights of the bounding boxes of X and Y.

Now, let an arbitrary 2D point $\mathbf{x}$ in a Euclidean space be defined as a homogeneous coordinate:

$$\mathbf{x} = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}, \tag{3}$$

where x and y are the horizontal and vertical components of the point. Then the matrix product $\mathbf{Tx}$ is an alibi transformation of the point $\mathbf{x}$ into a normalized space that enables an invariant comparison between X and Y in a Euclidean coordinate system.

The spatial distance between two points is then:

$$d(\mathbf{x}, \mathbf{y}) = \left\| \mathbf{Tx} - \mathbf{y} \right\|_2. \tag{4}$$

where $\| \cdot \|_2$ is the Euclidean (L_2) norm.

A total distance between X and Y is:

$$D(X, Y) = \frac{1}{n} \sum_{i=1}^{n} \left(d(\mathbf{x}_i, \mathbf{y}_i) w(i) \right), \tag{5}$$

where $n = |X| = |Y|$ is the (equal) number of sampling points in X and Y, and $w(i)$ is a weighting function that weights the contribution of each point-to-point distance. To treat all point-to-point distances equally, set $w(i) = 1$ for any i. Equation 5 means that two patterns that are identical have zero distance, and the more the sum of the point-to-point distances between X and Y increases, the more dissimilar $D(X,Y)$ becomes.

Finally, the template Y^* that best matches X is then found by comparing X against all Y_i in the gesture set Ω:

$$Y_i^* = \arg\min_{Y_i \in \Omega} D(X, Y_i) \qquad (6)$$

Space Time Cube Visualization
While equations 1-6 probably do not describe an optimal gesture recognition algorithm, the algorithm is remarkable simple and lends itself to a direct visual representation of the algorithm's behavior in space time cube representation.

This is achieved by mapping the points onto a plane in the space time cube, and letting the third axis represent time (or order) of the points.

The total distance $D(X,Y)$ between X and Y is then demonstrated to users by illustrating the corresponding point-to-point distances as connected bars. If the weighting function $w(i)$ is not constant in equation 5, the actual weighting for each point-to-point distance is revealed by coloring each connected bar using a pseudo-color scale.

Since the recognition algorithm has a direct visual analogy I am interested in studying whether completely accurately visualizing the behavior of the algorithm, affects users' perception and understanding of the actual implementation. This research question would provide a partial answer on whether it is worth effort to visualize more abstract statistical classification algorithms in an analogous manner, perhaps by transforming statistical feature comparisons into the concrete input space (e.g. the touch-display) that sampled the original input gesture in the first place.

I am also considering multiple other variations and extensions on how to visualize the above finger/pen-gesture recognition algorithm using space time cube.

Acknowledgements
The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement number 220793.

References
[1] Hägerstrand, T. What about people in regional science? *Papers in Regional Science 24*, 1 (1970), 7-24.

[2] Kristensson, P.O., Dahlbäck, N., Anundi, D., Björnstad, M., Gillberg, H., Haraldsson, J., Mårtensson, I., Nordvall, M. and Ståhl, J. An evaluation of space time cube visualization of spatiotemporal data. *IEEE Transactions on Visualization and Computer Graphics* (2009): in press.

[3] Kristensson, P.O. *Discrete and Continuous Shape Writing for Text Entry and Control*. Doctoral dissertation, Linköping University, Sweden (2007).

[4] Kristensson, P.O. and Zhai, S. SHARK^2: A large vocabulary shorthand writing system for pen-based computers. *Proc. UIST 2004,* ACM Press (2004), 43-52.

[5] Long, A.C., Landay, J.A. and Rowe, L.A. Implications for a gesture design tool. *Proc. CHI 1999*, ACM Press (1999), 40-47.

Temporal Data and Data Temporality: Time is change, not only order.

Catherine Letondal
letondal@lri.fr

INRIA Saclay
F-91405, Orsay, France

Aurélien Tabard
aurelien@tabard.fr

LRI - Univ. Paris-Sud & CNRS
F-91405, Orsay, France

Wendy Mackay
mackay@lri.fr

TELECOM ParisTech – CNRS LTCI
F-75013, Paris, France

Abstract
In this paper we describe two approaches to temporal data and how taking both into account can help improve the design of temporal data-based tools. *Time as order* considers temporal data as data that can be described by a time attribute, which can improve navigation or organization. *Time as change* considers temporal data as data that evolve over time.

Keywords
Time, temporal data, reflection, laboratory notebooks.

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

Introduction
The majority of tools supporting temporal data use time as a convenient property to help search data, or to get a global view on a series of information items. As such, time helps supporting navigation and organization tasks. In this paper, we would like to emphasize a different aspect, namely the fact that data itself evolve over time. This corresponds to a classical and quite well-established distinction in the history of ideas: time as order, and time as change [Wolff 2004]. The former is usually the physicists' view of time, as an objective

non-reversible arrow, where time is the fact that everything happen in a global ordering. The latter belongs to a more philosophical approach, where the focus is on the experience of time, mostly through the observation that changes, such as aging or processing, occur.

Using time attributes associated to data for navigation or organization purposes rely on the time as order paradigm. Focusing on processes and data changes takes the time as change approach. Temporal data can thus be described either as data having time attributes, or data needing a dynamic description, acknowledging data temporal dimension - what we call temporality.

Data temporality in the context of biology research

The work we present is grounded in our observation of biologists' relationship to information. We focus here on the way they organize information through their laboratory notebooks and various other tools, move it from one medium to another, and revisit it. We illustrate how the *time as change* paradigm helps to better capture the dynamic aspects of their data through three types of data temporality: TODO lists, reflective activity (biology research) and project management.

TODO items are a familiar type of temporal and dynamic data. Indeed, a TODO is an item that follows a typical path: to be done, with or without a deadline, then maybe imminent when hitting the deadline, then done. After being done, a TODO item can then be either thrown, or archived.

Scientific ideas also follow a path from their beginning through various steps that may include blackboard diagrams, discussions and meetings notes, experimental data, peer scientific publications, etc… This is related to the fact that biology researchers are in a creative process and **reflect** on their decisions in order to explore new leads or justify their decisions. Paper laboratory notebooks show this temporality of thoughts.

At a higher level, with their roadmaps and deadlines, **projects** also hold a temporal dimension. Yet projects and time organizations are often in conflict. Both provide dimensions to discuss activity but are they difficult to integrate [Tabard & al, 2007]. As researchers explore different projects, it is difficult to put order in their ideas, their framing changes and transitions or regrouping occur [figure 1] (*time as change*). When the projects are over, defined, researchers can refer to them as a whole and situate them in time (*time as order*).

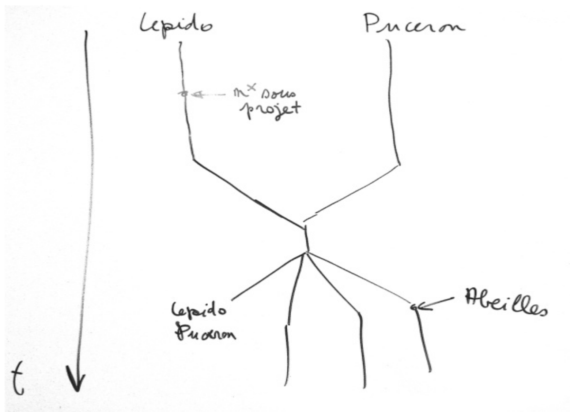

Figure 1: evolution of genomic projects organization over time. Common analysis to both projects are joined in one project, whereas other activities are spiltted.

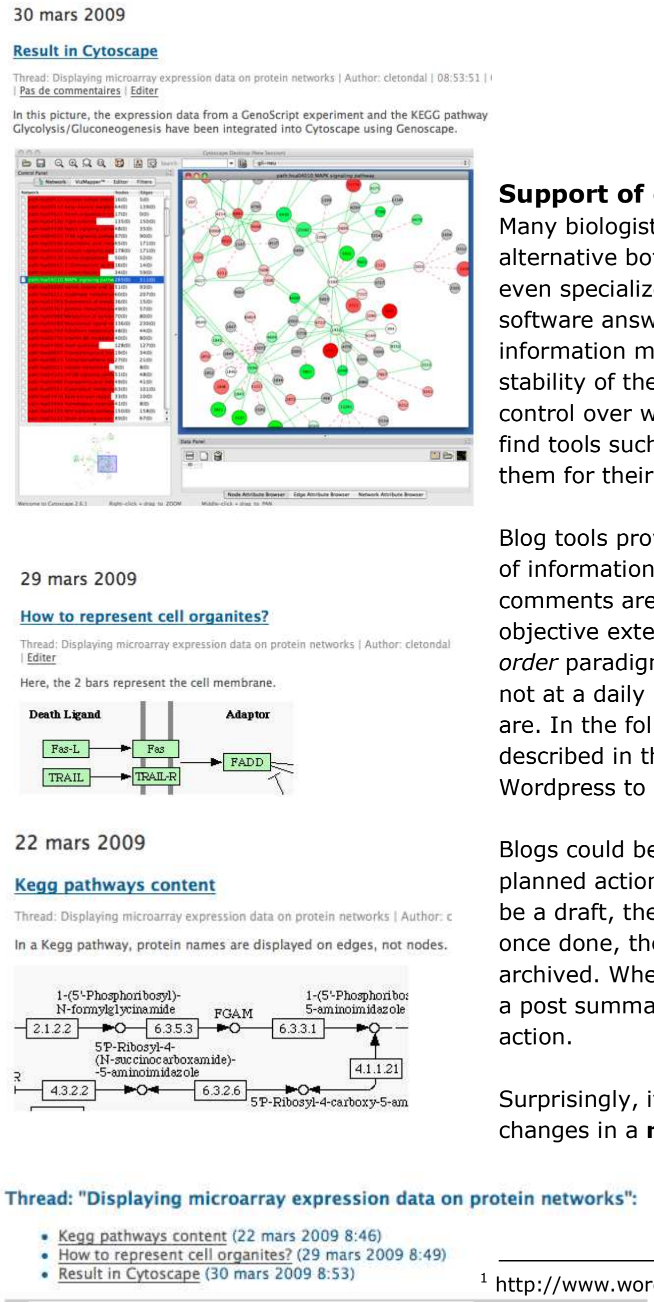

Figure 2: Thread of the evolution of a bioinformatics experiment [Clément-Ziza & al. 2009]

Support of data temporality in blogs

Many biologists are exploring blogging tools as an alternative both to paper laboratory notebooks and even specialized electronic notebooks. Blogging software answer three of their concerns regarding information management: durability of the data, stability of the system, confidentiality and privacy control over what they wrote. Furthermore, biologists find tools such as Wordpress[1] flexible enough to adapt them for their own need.

Blog tools provide a strong support for the chronology of information. Each post together with edits and comments are time-stamped, date becomes an objective external attribute. Blogs thus fit the *time as order* paradigm. Yet, they are organized at a post level, not at a daily level by default like laboratory notebooks are. In the following, we revisit the 3 types of data described in the previous part and see how we adapted Wordpress to better suit laboratory work.

Blogs could better support and integrate ongoing and planned actions. In the blog model, a **TODO item** can be a draft, then an email if a deadline is met. Then, once done, the item can either be thrown away or archived. When appropriate, the researcher may create a post summarizing the outcome of the corresponding action.

Surprisingly, it is quite difficult to express evolution and changes in a **reflective activity** over several posts: follow-ups of a post are available only through comments or trackbacks, which are second order writing - they do not appear as proper items in the blog index. Wordpress provides versioning, which enables to manage items that need to be changed. However, updating a post and just having access to its changes history does not capture the process over time, within the global stream of posts. Blog tools are designed as publishing tools; they do not support iterative thinking the way paper notebooks do.

Threading seems more appropriate to express data changes and processes. For example the evolution from raw data, to filtered data, to preliminary analysis of sample, to statistical analysis of large scale experiment. Everything refers to the same experiment, the data is the same, yet the information is changing. A biologist reporting an experiment lasting over several days or even weeks can keep track of the same "item" - the experiment - through following posts. Figure 2 shows a thread of 3 posts related to the same bioinformatics experiment. Threads are different from the idea of projects or categories. Threads are not defined in advance it is the information with the thread that define them. On the contrary **projects** or categories define the content they hold.

Medium temporality

Contrary to paper notes, computer files do not display the traces or versions that led to their final state. The transparent iterative edition capabilities of computer files hide their history to the users. Because of computers' editing flexibility, the transitions from action items, to reflection, to finished and articulated project, are lost.

It is thus important to remember that the medium and the information interact. The medium influence what

[1] http://www.wordpress.org

biologist will save, while biologists pick the medium that fits best their needs. Paper scraps suit short notes, whereas biologists use computer tools to document their digital activity.

Paper holds temporal properties which are not yet integrated in computer. Paper notes and notebooks display their age, the number of iterations they went through. Furthermore, users can program paper decay: rather than pasting a sheet in a notebook, leaving it loose hints that it will disappear at some point in time.

As more activity happens online, biologists use less paper notebooks. However they feel they "lose discipline" when moving to complete digital notebooks. Paper allows smooth transitions between action items to reflective activity and project management. When saving a short note that may have no value later, paper scraps are still more efficient. Paper notebooks support easy re-visitation and reflection, by providing a constrained writing, with a limited space (only shorts edits are possible) and chronological order.

Compared to paper artifacts such as laboratory notebooks, computer files do not offer a proper structure to manage temporal and evolutive data. Being editable, it is difficult to reach a definitive state of a computer-based document. This is one of the reasons why biologists still use paper notebook for their linear and constraining structure.

Computer-based notebooks should provide a similar temporal structure, as blogging tools do with their automatic time-stamping of posts. Even in the presence of discordance between an automatic timestamp and the actual date of an experiment, we observed that biologists prefer to keep this automatic and objective date, rather than modify it - and manually add the *real* date of the experiment.

Conclusion

Providing tools to navigate and search among time-stamped data is not enough. There is also a need for mechanisms to support transformations and processes over time, both for scientific data and scientific ideas. These mechanisms should not only help the user visualize but also *express* time and change.

Citations

[1] Clément-Ziza, C., Malabat, C., Weber, C., Moszer, I., Aittokallio, T., Letondal, C. and Rousseau, S., 2009 Genoscape: a Cytoscape plug-in to automate the retrieval and integration of gene expression data and molecular networks. Bioinformatics, Oxford University Press, submitted

[2] Wolff, Francis. Dire le monde, Paris, P.U.F., 1997, 220 p. 2ème édition complétée avec une Préface inédite et deux annexes, P.U.F., « Quadrige », 2004, 284 p.

[3] Tabard, A., Mackay, W. E., and Eastmond, E. 2008. From individual to collaborative: the evolution of prism, a hybrid laboratory notebook. In *Proceedings of the ACM 2008 Conference on Computer Supported Cooperative Work.* CSCW '08. ACM, 569-578.

Annotating Probabilistic Semantics for Time Expressions

Gregory Marton and Piotr Mitros
MIT CSAIL
gremio@mit.edu

Temporal references pervade our communication, time is primitive to computing, and yet the semantics of time expressions are computationally not well represented. We propose a simple probability-centric model for time semantics, and we have created a user interface intended to elicit data within this model. We hope that the interface will be usable by untrained annotators, through the Amazon Mechanical Turk harness, and we are looking for feedback on ways to improve that user experience, and thus the quality of the data we can gather.

Introduction

Meanings for time expressions are most clearly annotated today by projects like TimeBank (Pustejovsky *et al.* 2002), and the related TIMEX2 and TIMEX3 annotations in the ACE Time Expression Recognition and Normalization (TERN) corpora (Ferro 2004). In this project, we plan to re-annotate a subset of the existing TIMEX annotations on the existing time-based corpora, essentially providing an alternative annotation for some of the ground-truth tags.

The alternative annotation will approximate probability distributions on a timeline, so that, from a reference time of 2009-01-01, the probability distribution for the expression "tomorrow" will put (at least the bulk of) the probability mass on 2009-01-02, and may give proportionally more probability to daytime hours than to nighttime ones, based on annotator preferences.

This paper presents a user interface intended to elicit information about the probability distributions that we suppose naive annotators to have. They are asked to look at individual time expressions in the context of the sentences in which they appear, and to mark when they think the events referenced most likely took place. The time expressions are drawn from an existing corpus of manually-annotated time expressions of the kinds that we are interested in.

Motivation and Scope

We plan to use two primary corpora that include TIMEX annotations. The ACE 2005 TERN task's associated corpus is the larger of our corpora, and includes 650 documents in a variety of genres including newswire, broadcast news, telephone conversations, blogs, and usenet. In these documents, there are 5469 TIMEX2 entities marked, of which 1933 are unique (about a third). The TimeBank corpus is smaller, containing 186 newswire documents, with 1416 TIMEX2 annotations, 405 unique. The long tail emphasizes the need to understand the semantics of such time expressions compositionally.

Among the great variety of temporal meanings expressed in language, we will focus our effort on the simplest kind of temporal location, most often represented by points or anchored intervals on a time line. We are interested in clearly delimited (if vague or ambiguous) time expressions like "three hours ago", "last Wednesday", and "early Tuesday". We explicitly exclude unanchored durations such as "three hours", repeated times such as "Wednesdays", and non-specific times such as "on a Tuesday". We also explicitly exclude time references by tense and aspect, and references to events such as "during the trial". The meanings of the excluded forms, we conjecture, may become derivable once a sufficient understanding of the included locative expressions is reached.

To concretize these exclusions, we refer to the six primary fields specified in the TIMEX annotation guidelines (Ferro *et al.*): the time value, the modifier if any, anchor value and direction if any, whether the time expression denotes a set, and whether it is non-specific. We are excluding non-specific times and sets, so any ground truth elements in which these fields have a value "YES" will not be considered. Periods or durations are marked with a "P" at the beginning of the value field, and these are also excluded. For simplicity, we will presently exclude any values that indicate a time zone.

Within the scope of simple locative time expressions, we were motivated to take a fresh look at the model of semantics because the TIMEX annotations seem difficult to use in further computation. We cannot summarize the near 50 pages of relevant guidelines here, but will show a few examples to illustrate the difficulties.

- The values "2009-01-08" and "2009-W2-4" (the fourth day of the second week of 2009) refer to the same date, though the first is preferred, where it is possible to specify. "last week" is annotated in the week format, whereas "last Thursday" is annotated in the date format. A "weekend" is a special day of the week, and "weekday" does not appear in the standard.
- Similarly to the "weekend" for days, seasons and times of day are specified as special months and special hours, respectively. This means that no human judgement is given as to when a "summer" event most likely took place, nor what "evening" likely means. To add to the confusion, "winter" and "last night" must choose a year and a day respectively to associate with—it is impossible to express that they might span our measurement boundaries.
- Annotators must *use only the head noun to determine the precision of a value.* The values for "three years and six months ago" and "three and a half years ago" are different, the first at month-level precision, the second at year-level precision, but with a decimal 3.5 years ago specified. What is the meaning of precision when decimals, undecorated with precision information, may be used? The year form perhaps includes more margin for error than the month form, but no such judgements are elicited from the annotators; the choice must be made based on the head noun alone. Worse, the precision of the head noun may be misleading: "earlier this month" by this rule gets annotated at month-level precision (VAL="2008-01" for a 2008-01-08 article), but we have much more information.
- The modifiers similarly constrain the annotator, eliciting too little information. "early this month" would have a MOD="EARLY" attribute, but the distinction between "over two years ago" and "just over two years ago" is lost, as is the distinction between "nearly 30 days from now" and "in a maximum of thirty days".
- Anchors specify inclusiveness of the anchor point, (AFTER vs. AS-OF, BEFORE vs. ENDING), but they do not allow the annotator to capture any other sense of the relation between the anchor and the target time. When we talk about "Ford models before the 2003", we naturally exclude years in the 1400s, when there weren't cars. More explicitly, "before tomorrow's 10am meeting" in contextual future tense cannot be before the speech act today. This problem shows up also in the FUTURE-REF annotations and kin: "in a few days" is marked with VAL=FUTURE-REF, but we have much more information.

Taking a probabilistic view allows annotators to capture nuances of meaning, and makes computation with the annotation results more straightforward. We exclude these expressions for the present, but the extension to multiple events that happen on "Wednesdays" and to probability distributions spread over a discontinuous range like "a rainy Tuesday in November" seem more-or-less natural. However, the probabilistic view of time expressions has some clear drawbacks, too:

- It is difficult conceptually to capture the notion of time zones. These could be viewed as parallel time lines, or as offsets to a particular time line, but we have not found a solution within the probability-centric model. At best, this may be maintained as external information, as in the TIMEX2 standard.
- Discontinuities such as the BCE/CE boundary and the Julian conversion can be viewed as a mapping from an idealized timeline onto various calendars. Like timezones, the particular calendar onto which one maps must be externally maintained.
- Fiscal years vary from company to company, and people quite freely make inferences about the fiscal years for a company without knowing the exact dates of its fiscal year boundaries. The TIMEX standard allows values like "FY1998" without further specification, but this kind of vagueness is not possible with a single time-line representation. This goes beyond externally maintaining a calendar, towards the kind of reasoning with uncertainty that representations for non-specific TIMEXes likely entail.

With these caveats, and despite those we are likely to discover in moving forward with this annotation format, we suspect that a probabilistic view of time expression semantics will make the most sense for many applications.

The Annotation Interface

We have implemented a preliminary time expression annotation interface, and we plan to integrate it with the Amazon Mechanical Turk harness to elicit annotations from naive users. A screen capture of one annotation is shown in Figure 1.

In the user interface, we show a single timeline at multiple levels of granularity. As one zooms in on the timeline, new levels of granularity become available. In the figure, minutes are visible, but seconds do not yet appear. As one zooms out, hour widths will get smaller,

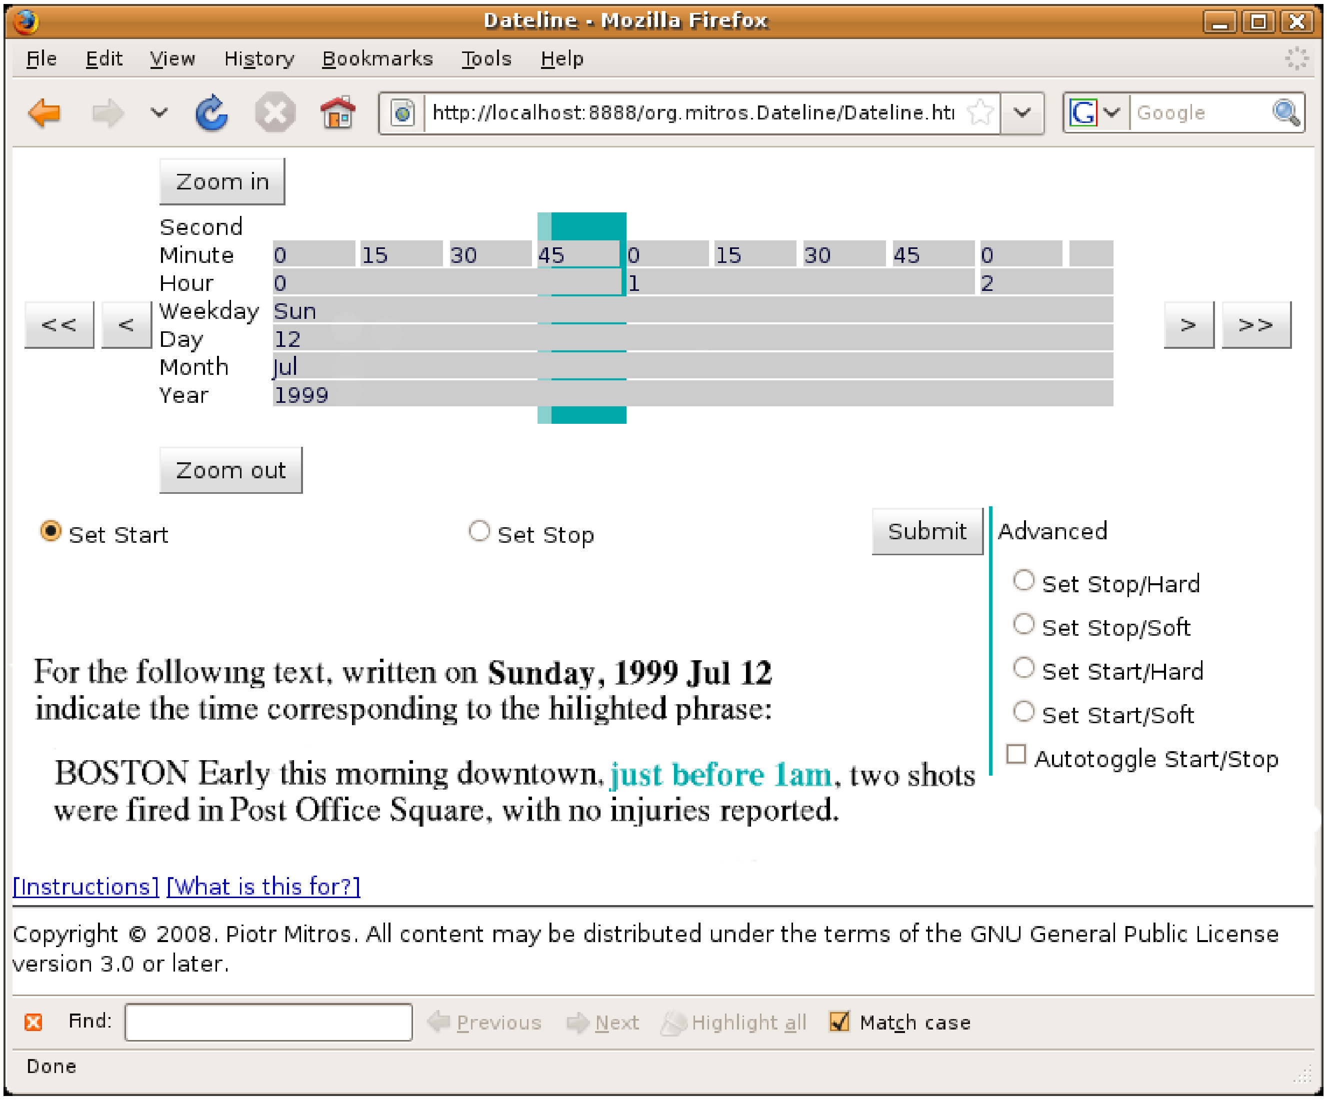

Figure 1: A typical example has been annotated, with the target interval's start and end times at the same granularity (143 min / 900 px), but with the interval start "soft" and the end "hard".

minutes will disappear, eventually the days (and weekday with them) would collapse, up to groups of tens, hundreds, thousands of years. In this case, both the starting and ending points were specified at the same level of granularity, and that will be recorded in terms of the number of minutes per pixel.

The user is given a prompt, specifying the article date, to which all time expressions are expected to refer, if no other reference is clear. We plan to show only one sentence for the prompt, as a tradeoff between giving too little context, where some reference times may be missing or wrong, and giving too much context, where the annotator would have to spend far too long on an individual annotation. If a user cannot determine the time meant, the form can be submitted without annotation.

The timeline display starts at roughly day-level granularity, centered on the article reference date. The user must zoom in and out, and travel by small or large increments forward or back in time using the arrows on the right and left, in order to bring an endpoint of the relevant interval into view. They may then set either the beginning or the end of the given interval, and may choose to use a sharp or a gradual boundary. Sharp boundaries expand to gradual ones as one zooms in on them — they take their granularity into account.

At present, the timeline does not implement leap days, nor the BCE/CE discontinuity (lack of year zero), and we limit the zoom level to years before 1800 to avoid problems with the Julian calendar (used until 1752 in England). This means that, for example, Isaac Newton's birthday "360 years ago today" (4 January 1643 in the modern calendar 25 December 1642 in the calendar of the time) cannot be accurately represented.

Contributions

We have presented a new probability-based model for time expression semantics, and we have shown a preliminary user interface for eliciting human annotations for such semantics from naive users.

The annotations are not meant to embody a single correct answer, but to provide information towards a probability model. Rather than measuring inter-annotator agreement, we will provide a probability distribution, with more weight on regions that more annotators have chosen to include. For this reason, multiple annotations for each expression will be key to the success of such a project. At $0.01 per annotation, 30 annotations per time expression in our corpus will cost about $2000. This is competitive pay on Amazon Mechanical Turk, and is relatively inexpensive for a corpus annotation project.

For the workshop discussion, we plan to demonstrate the user interface, to ask participants to annotate some data, and to ask for feedback on how we might improve the user experience and measure the resulting data quality.

References

Ferro, L.; Gerber, L.; Mani, I.; Sundheim, B.; and Wilson, G. TIDES 2005 Standard for the Annotation of Temporal Expressions.

Ferro, L. 2004. Tern evaluation task overview and corpus. In *Time Expression Recognition and Normalization Workshop.*

Pustejovsky, J.; Belanger, L.; Castaño, J.; Gaizauskas, R.; Hanks, P.; Ingria, B.; Katz, G.; Radev, D.; Rumshisky, A.; Sanfilippo, A.; Saurí, R.; Setzer, A.; Sundheim, B.; and Verhagen, M. 2002. Nrrc summer workshop on temporal and event recognition for question answering systems.

Temporal tribulations: the trouble with time selection

Peter McLachlan
University of British Columbia
Vancouver, BC V6T 1Z4 Canada
spark343@cs.ubc.ca

Tamara Munzner
University of British Columbia
Vancouver, BC V6T 1Z4 Canada
tmm@cs.ubc.ca

Stephen North
AT&T Labs Inc. - Research
Florham Park, NJ 07932 USA
north@research.att.com

CHI 2009, April 4 – April 9, 2009, Boston, MA, USA
ACM 978-1-60558-247-4/08/04.

Abstract
Even within a single application domain and data set relevant time intervals can scale from milliseconds to years. One example of this is system management data used in the planning and operation of data warehouses. We propose a generalized modification to a visualization method introduced by Stolte [1] for scalable, intuitive time range selection. This method would allow the user to see the entire available data range and visually select down to an arbitrary precision. We would like to investigate other interaction methods that may support this type of scalable navigation while preserving visibility of relevant time ranges.

Keywords
visualization, time-series, hci, interaction

ACM Classification Keywords
H.5.m Miscellaneous

Introduction
Current desktop computer systems execute a primitive instruction in two to three nanoseconds. Data collection for debuggers can take place on the scale of microseconds. Operating system logs are collected on the order of milliseconds. Network monitoring logs are

collected on the order of minutes, hours, days and months. Time is a sufficiently important data dimension for computer system administrative tasks to warrant the development of an entire protocol and toolset for synchronizing time across computer systems. The Network Time Protocol is used globally in the enterprise by desktops, servers, mainframes and networking equipment to accurately synchronize time with authoritative time sources.

Why is accurate time so vital to network computer administration? Primarily it is used for the purposes of audit. Log records are kept of system transactions and events from the level of individual applications all the way to the lowest level of the operating system kernel operations. These log files are used to understand system state, perform forensics during failure or intrusion and also during capacity planning. When analyzing a single system it is preferable, but not critical, to have time synchronization so that external events can be compared to internal computer events. However, when analyzing data from networks of interrelated systems accurate time synchronization becomes critical in problem determination. Having systems with unsynchronized time will make finding temporal correlations impossible.

Analysis of network management data can take place at many different intervals of time. For some analysis data intervals may be specified in nanoseconds. In other cases, the smallest sample size can be a time interval of an entire day or month. We would like users to be able to easily and interactively locate and define ranges within these macro/micro scales.

Network Management Time-series visualization research

We are developing a visualization system for massive scale time-series data called LiveRAC [1]. Our system enables users to display data at multiple levels of detail. It is intended to support both very fine-grained time-series on the scale of milliseconds all the way up to time-series that span multiple years in duration. Interacting with this data has been accomplished to date using a double-edged slider that changes scale based on the size of the time range selected. This approach has a number of drawbacks including difficulty with precision when interacting with large time scales and the ease of overshooting targets. Our experience with our implementation in the field suggests it is worth investigating alternative approaches.

Large dynamic time range challenges

We define large dynamic time range as any time range that scales seven orders of magnitude. There are two key challenges to effectively supporting large dynamic time ranges, described below.

Variable time intervals

At the macro scale, we are used to working with a variety of time intervals that are based off of a solar calendar. These units include minutes, hours, days, weeks and months. These are the units that have the most meaning to humans and cannot be omitted. Any time range selection tool must support display of macro time scales using these and other non-Gregorian derived intervals.

Time is (practically) unlimited and continuous
Theoretical physics aside, time continues long enough into the past and far enough into the future that it's unlikely that the user is interested in all of it. Furthermore, time can theoretically be measurably subdivided down to a Planck time (5.4×10^{-44}s) interval. Any solution to our interaction problem requires that we only display relevant scales on which we have both data and resolution.

The requirement that we only show some chunk of time means there are now two ranges involved in our display. The first is range total, R_t and the other is range selected, R_s where $R_s \subset R_t$. If each range is defined by the user by means of a beginning and end time, the user must define a total of four times. Furthermore, we must prevent the user from committing errors such as making R_s larger than R_t or from shrinking R_t to be smaller than R_s. Defining four times is cumbersome – an automated approach for finding the total range would be strongly preferred, restricting the user to picking only a single relevant range.

Proposed solution

We believe that a visualization approach used in a paper by Stolte [1] may offer a framework that can be generalized to create one solution for resolving scalable time range selection. The approach by Stolte shown in Figure 1 consists of a set of stacked strips representing different time intervals. In this approach all selected intervals use a quantum of milliseconds. Each strip in the visualization stack represents a smaller area of time. The implementation in Stolte's paper only supports three levels of strip chart.

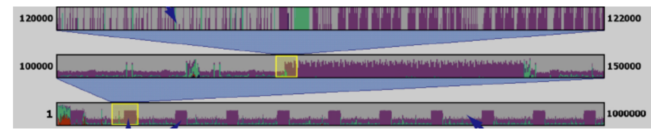

Figure 1 – Stolte's approach [1] visualizes data over time at different intervals. We propose a similar mechanism as a generic time selection widget.

We propose a solution inspired by Stolte, implementing a flexible number of strip charts depending on the level of resolution required for the data. Each strip uses a different quantum of time interval. A generic widget can be constructed that takes as parameters the different levels of resolution required. These resolution levels can include intervals defined by Gregorian or other calendar systems. We have mocked up an example of such a system in Figures 2 - 5.

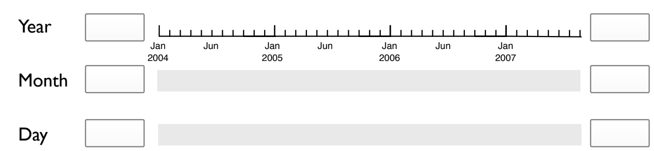

Figure 2 – Years reflect data edges, no user selection.

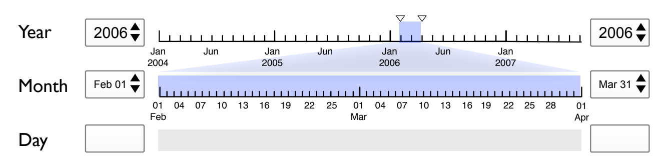

Figure 3 – User selects a range from the year strip.

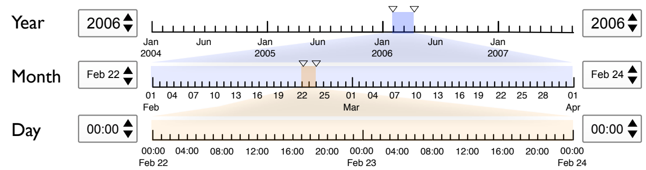

Figure 4 – User selects a range from the month strip.

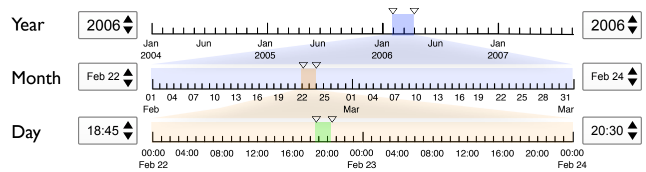

Figure 5 – User selects a range from the day strip.

The widget takes as inputs the total range of the available data, and the desired user intervals. It requires objects that can translate between the target time intervals. These objects may be simple wrappers for pre-existing calendar and date objects.

Discussion and Conclusion

Our approach should address some of the most significant problems of interacting across multiple scales of time. However, it leaves open questions. What if the user is viewing live data that changes? Do we modify the uppermost strip on the fly to indicate the arrival of new data? Also, this approach consumes a considerable amount of screen real estate. Some commonly used GUI techniques could help address this although there are tradeoffs. The widget could support folding unneeded strips or placing the strip charts in a scrollable region. Both of these involve hiding potentially important components of the widget from the user and increasing interaction time. We believe there may be other innovative approaches that outperform the one described in this paper, or suggestions on improving this approach. We look forward to discussing how other researchers have coped with large-scale dynamic time range selection.

Bibliography

[1] McLachlan, P., Munzner, T., Koutsofios, E., North, S. LiveRAC: interactive visual exploration of system management time-series data. *Proc. ACM Conference on Human Factors in Computing Systems (CHI), 2008, 1483-1492.*

[2] Stolte, C. Bosch, R., Hanrahan, P., Rosenblum, M. Visualizing Application Behavior on Superscalar Processors. *Proc. IEEE Symposium on Information Visualization (InfoVis)*, 1999, 10-17.

Temporal Data in a Health Self-Management Application

Yevgeniy Medynskiy
GVU Center
Georgia Institute of Technology
eugenem@gatech.edu

Andrew Miller
GVU Center
Georgia Institute of Technology
andrew.miller@gatech.edu

Jae Wook Yoo
GVU Center
Georgia Institute of Technology
jyoo@gatech.edu

Elizabeth Mynatt
GVU Center
Georgia Institute of Technology
mynatt@cc.gatech.edu

CHI 2009, April 4 – 9, 2009, Boston, MA, USA
ACM 978-1-60558-246-7/09/04.

Abstract
Individuals frequently take an active role in managing day-to-day aspects of their health, like improving nutrition or increasing physical activity. Clinicians also increasingly teach health self-management skills to patients with a range of chronic illnesses, such as diabetes, hypertension or arthritis. In this position paper, we present our initial work in designing and developing *Salud!*, a web-based platform for supporting health self-management. *Salud!* will allow its users to track personally-relevant aspects of their everyday life, and provide visualization and analytics tools with which to make sense of the resulting datasets. In effect, *Salud!* is a health-oriented, capture and analysis tool for temporal data. We describe the features of *Salud!* that will enable users to easily capture temporal data, and to use this data in a number of ways. We conclude by discussing how we have structured *Salud!*'s data storage system, and our plan for addressing the challenge of designing temporal data visualization and analytics tools for a broad, lay user base.

Keywords
Health self-management, temporal data, analytics

ACM Classification Keywords

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

Introduction

Patients with a chronic illness, such as diabetes, heart disease or asthma, are often called upon by clinicians to actively participate in the day-to-day management of their health. Health self-management, which may include making healthy decisions about diet or exercise, quitting unhealthy habits and solving other health problems, complements traditional healthcare practices and supports patients in living the best possible quality of life with their condition. Studies have found that successful health self-management education and group support programs improve patient outcomes and reduce healthcare costs [1,5].

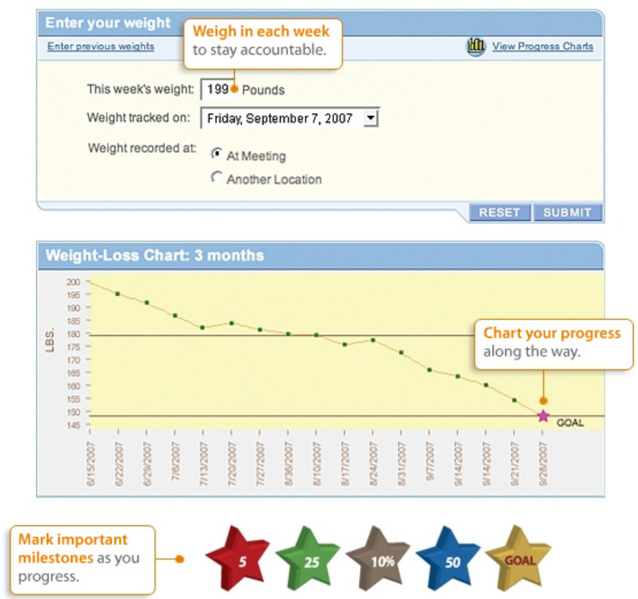

Figure 1. WeightWatchers e-Tools allow users to chart their progress toward a weight goal.

A feeling of self-efficacy—confidence to carry out a behavior necessary to reach a desired goal—is central to effective health self-management. It is enhanced when patients succeed in solving problems which they themselves identify. Improving self-efficacy and motivation is also important for individuals without chronic illness, but who are trying to lose weight, eat better or otherwise promote their health and wellness.

A growing number of commercial and research systems attempt to support individuals in achieving their health goals. For example, Fish'n'Steps [4] and Chick Clique [7] leverage competition and peer-pressure to motivate users to become more physically active. Consumer technologies, such as iPod Nike+, potentially support motivation and increase self-efficacy by allowing individuals to track their progress in achieving a specific goal. Similarly, many online services such as WeightWatchers e-Tools (Figure 1) allow individuals to record and track a wide range of information (often temporal) relating to health and wellness [3]. The features which these services provide have only limited support for health self-management, however. Most provide only basic views of a user's data, which may be laid out in a table or calendar, or plotted as a line through time. While this may help users stay motivated by visualizing long-term trends, there is little support for decision-making or identifying and overcoming obstacles to improvement. Additionally, users are limited to the particular measurements a service supports (e.g., weight), which may be insufficient for a range of health goals.

To begin addressing some of these limitations, we are designing and developing *Salud!*, an open, web-based platform for personal health self-management. *Salud!* has three main design objectives: support users in capturing a wide range of data about their everyday experience; allow users to visualize this data and provide a repertoire of easy-to-use analytics tools for identifying patterns and trends; and structure problem-solving activities which lead to achievement of users' health goals. The goal of our research is to learn if, and how, such a tool can help users better understand the factors which affect their health, and support them in attaining self-identified goals.

The Salud! Platform

Users' interactions with *Salud!* are centered around Logbook objects. A Logbook stores a temporal record of some measurement or variable which the user is tracking. For example, a user may create a Logbook to record information about all of her meals and snacks during the day. Each entry in the Logbook would record the time of the meal or snack, and any pre-defined annotations, such as a photo of the meal, its approximate calorie count and/or a list of its primary contents. Other Logbooks could be used to record the users' blood pressure, self-reported stress level, and other types of data. While *Salud!* will be pre-populated with Logbooks for common health and wellness metrics which users can being using immediately, users will be able to create new Logbooks to track other variables which are important or interesting to them.

Salud! will also provide users with easy-to-use tools with which to visualize their data and perform simple data analysis. Users will be able to track trends over time, explore relationships between different variables, and manipulate the data in other ways. Because all of the data is timestamped, we are considering a range of temporal visualizations with which to provide views of the data, in order to allow users to answer a variety of questions. We are also building a number of statistical tools into *Salud!,* which would allow users to smooth noisy data, compute correlations, and perform other kinds of simple analysis on the data.

The final element of *Salud!* will be a set of guides and decisions-support mechanisms which will support users in defining and reaching concrete health goals. Depending on how *Salud!* will be deployed, these tools may be designed to be used by individuals themselves, or with support from a clinician or health educator. A detailed discussion of these tools is out of the scope of this position paper. In the remainder of this section, we will more thoroughly describe the data capture and analytics tools that we are building into *Salud!*.

Capturing Temporal Data

Capturing pertinent health and wellness measurements is a key aspect of effective health self-management. Individuals managing a chronic illness often keep written diaries or simple computer records of the behaviors and measurements which are significant to their condition [6]. Diabetics are encouraged to monitor their diets and blood glucose readings, individuals trying to lose weight may log their exercise routine as well as their weight measurements, etc. However, aside from the obvious issue of motivation, manual data capture is often difficult because individuals need to remember the specifics of an event (or even that one occurred), until it can be recorded. Because of the effort and overhead involved, individuals rarely track more than a few variables.

11 Dec 2008 8:35 AM

Calories: 250

Contents
Fried egg
Fruit Salad
Coffee

11 Dec 2008 12:22 PM

Calories: 700

Contents
Spinach
White rice
Steak
Diet Coke

11 Dec 2008 6:49 PM

Calories: 800

Contents
Bread
Roast beef
Veggies
Water

Figure 2. Sample data which may be recorded in a Logbook.

We are building a number of different data capture methods into *Salud!*, which we hope will significantly ease this task, making it possible for individuals to keep more consistent and more detailed records of a greater number of variables. In addition to a web-based interface, *Salud!* will accept new data records via email, instant messenger, and SMS/MMS messages. This will allow users to record an event (e.g., a meal, onset of pain, etc.) at the time of its occurrence, or log a measurement (e.g., blood pressure) immediately after taking it. Additionally, *Salud!* will allow users to create reminder schedules for measurements they would like to record regularly. The system will then send them reminders via a specified communication medium, and the replies to these reminders will be parsed and added to the appropriate Logbook.

Consider a hypothetical user, Sareen, who would like to change her eating habits so as to cut calories, while avoiding fatigue (i.e., maintaining her "energy levels") between meals. Sareen wants to track the calories and contents of her meals and snacks in one Logbook, and routinely self-report how fatigued she feels, on a scale from 1 to 5, in another Logbook. She can collect meal data by photographing her meals and snacks with her camera phone and immediately sending the picture to *Salud!*. The photo would serve as a placeholder for the meal, and record its time. Every evening (or every several days), Sareen can then log in to *Salud!*'s web interface and annotate recently added photos with their contents and calorie counts. To collect data about her energy levels throughout the day, she can create a reminder schedule which will send a short question to her work email address every weekday, at 10:30 AM and 4:00 AM. By replying to that message with '1', '2', '3', '4' or '5', Sareen can easily track this variable.

Because we cannot foresee all of the types of data which users may want to collect, or how they may find it convenient to do so, we have created a simple REST API through which data entry occurs. The email, IM and SMS/MMS functionality is implemented as independent services that interact with this API. Other services, which we may implement in response to user demand, or which may be created by other researchers or technically-oriented *Salud!* users, could allow data streams to be captured by in-home sensors (e.g. "smart" scales, glucometers or sphygmomanometers) or be imported from other online data sources (e.g. another website, Facebook or even a weather service).

Visualization and Analytics Tools for Temporal Data

The second main component of the *Salud!* platform is a visualization and analytics system. This system will allow for open-ended exploration of captured data, and should also support users in addressing specific questions or problems regarding their health. We expect that having the ability to view and analyze a history of their data over time will make it easier for users to manage their health by establishing goals, monitoring progress and solving problems which stand between them and their goals.

Most of the data visualizations we are currently planning to include in *Salud!* will provide views of the users' data on a timeline. However, a flexible UI is needed to enable users to address different questions and concerns. Consider Sareen, our hypothetical user from the previous section, who is interested in understanding the relationship between her diet and fatigue. She may start her exploration of by plotting several weeks' worth of energy level self-measurements on a timeline, and explore various theories by superimposing meal-related data over this graph. By adding a dot to the timeline for every meal or snack she has, she can explore whether she experiences more fatigue if she skips breakfast or lunch. Later, she might want to see the meals' calories plotted as well, and highlight times when she had eaten specific food items ("Do I have more energy when I eat a low-calorie meal that includes dairy?"). To see if her weekly calorie intake has been dropping, however, Sareen would need a different view—perhaps a bar graph, with total calories bucketed by week.

Simple analytics tools will also be available to Salud! users. For example, applying a moving average line or a linear regression to a plot could help smooth out the noise in day-to-day changes of certain variables (e.g. weight), and better display overall trends. If presented appropriately, statistical tools like correlation coefficients may also help users think and act more confidently. Finding the right set of tools to meet users' needs, and providing access to these tools in an effective, usable way is the main challenge we foresee in the development of *Salud!*. In the next section, we will describe our plan for meeting this challenge.

Key Challenges

In this section, we describe two design problems which we have considered in our work with *Salud!*. First, we describe how we have structured the temporal data store. We then discuss our plan for iteratively designing *Salud!*'s visualization and analytics tools.

Structuring the Personal Health Dataset

Salud!'s data storage system is designed to allow timestamped data to be submitted and edited through a variety of interfaces, and then to make this data available to visualization and simple analytics services. From a user's perspective, data entry and data analysis involve two different ways of working with temporal data. For input and navigation, grouping multiple data fields in a Logbook—e.g. "meal photo," meal "contents list," and "calories"—helps the user easily track and review semantically related data. However, during the analysis process, it is usually necessary to isolate specific data fields, or manipulate multiple data fields from different Logbooks.

To allow for both types of modes of use, we structured our data store in the following manner. A user's dataset consists of one timestamped data stream for each individual data field they track. Each data field can be associated with at most one Logbook—in the example above, "meal photo", "contents list" and "calories" are each data fields associated with one Logbook ("Meals & Snacks"). Users then capture data at the Logbook level. When a user creates a data record, the system records a timestamped entry in a specified Logbook, and also adds timestamped entries to those associated fields which the user included in the data record. For example, sending a photo of a meal to *Salud!* would create a new entry in the "Meals & Snacks" Logbook, and a new entry in the "picture" field. Later, when the user adds the meal's calories and contents, new entries would be created in the "calories" and "meal contents" fields, with the timestamps matching the timestamp of the Logbook entry. The timestams of Logbook entries thus serve as a kind of primary key, making it easy to iterate over all entries in a Logbook and to look up the data fields in each entry. Because individual data fields are also timestamped, they can be isolated easily as well—even if a data field doesn't have a value for each Logbook entry.

This data storage mechanism results in some amount of redundancy (timestamps are stored multiple times), but has the benefit of making UI development more straight-forward. By associating multiple fields, Logbooks allow individual users to structure data collection in a way that is logical to them. Also, because each Logbook consists of data streams which may be accessed independently, it is straightforward to create analytics tools that allow users to visualize and manipulate multiple data fields from different Logbooks.

Designing the Visualization and Analytics Tools
The most important challenge we face going forward is designing *Salud!*'s visualization and analytics tools. While we hope to provide our users with a robust and relatively powerful system for understanding and making decisions about their health, we are keenly aware of the need to keep these tools intelligible and easy-to-use for a lay audience.

We will need to understand how users think about the trends and changes in their health in order to design useful tools. In particular, the ways in which users conceptualize data—in terms of frequency, durations, aggregation, etc.—will dictate the types of data views and visualizations that will be most important to include. The meanings which users will want to extract from the data, and the types of questions they will want to ask will similarly provide guidance for *Salud!*'s analytics tools. We will need to strike a balance with the number of features and options, how structured interactions are, and even the terminology and jargon in the interface, to provide an application that is sufficiently open-ended, without being overwhelming or unnecessarily technical.

We plan to develop these features of *Salud!* iteratively, working closely with a group of 12-15 early adopters who are specifically motivated and interested in using such a system. We will provide these individuals with a prototype version of *Salud!* and examine how they begin to appropriate its features into their health management strategies. We will iterate on the design and functionality of the system by rapidly acting on their feedback and feature requests. In this way, we expect to converge on a version of *Salud!* ready for more widespread deployment with a broader user base.

References

1. Bodenheimer, T., Lorig, K., Holman, H., and Grumbach, K. Patient Self-management of Chronic Disease in Primary Care. *JAMA 288*, 19 (2002), 2469-2475.
2. Hesse, M. Bytes of Life. *The Washington Post*, 2008, C01. http://www.washingtonpost.com/wp-dyn/content/article/2008/09/08/AR2008090802681.html
3. Lin, J., Mamykina, L., Lindtner, S., Delajoux, G., and Strub, H. Fish'n'Steps: Encouraging Physical Activity with an Interactive Computer Game. In *UbiComp 2006: Ubiquitous Computing*. 2006, 261-278.
4. Lorig, K.R., Sobel, D.S., Ritter, P.L., Laurent, D., and Hobbs, M. Effect of a self-management program on patients with chronic disease. *Effective Clinical Practice 4*, 6, 256-62.
5. Mamykina, L., Mynatt, E.D., and Kaufman, D.R. Investigating Health Management Practices of Individuals with Diabetes. *Proc. CHI 2006.* ACM Press (2006), 927-936.
6. Toscos, T., Faber, A., An, S., and Gandhi, M.P. Chick clique: persuasive technology to motivate teenage girls to exercise. *Ext. abs. CHI 2006.* ACM (2006), 1873-1878.

Author Bios

Yevgeniy "Eugene" Medynskiy is a third-year Ph.D. student in the Human-Centered Computing program at the Georgia Institute of Technology. His research focuses on developing technology to help individuals manage day-to-day decisions that affect their health and wellness.

Andrew Miller is a first-year Ph.D. student in the Human-Centered Computing program at the Georgia Institute of Technology. He also holds an M.S. in HCI from Georgia Tech and most recently worked as a User Experience Designer at Schematic.

Jae Wook Yoo is a first-year M.S. HCI student at the Georgia Institute of Technology. He most recently worked as a mobile phone software engineer at Pantech.

Elizabeth Mynatt is the Director of the GVU Center and a Professor in the College of Computing at the Georgia Institute of Technology. She also directs the Everyday Computing Lab, whose members examine the human-computer interface implications of having computation continuously present in many aspects of everyday life.

Temporal Views on Desktop Activity

Brian R. Meyers
Microsoft Research
One Microsoft Way
Redmond, WA 98052 USA
brianme@microsoft.com

CHI 2009, April 4 – 9, 2009, Boston, MA, USA
ACM 978-1-60558-246-7/09/04.

Abstract
For several years, we have been building views of personal interactions with desktop computers across time. Each of these views offered something interesting but none of them succeeded completely. We describe the various approaches and highlight some of the challenges each faced.

Keywords
Temporal data, Desktop activity management

Introduction
We are interested in both analyzing and providing a visualization of the user's interaction with their desktop computer. Our earliest attempts only dealt with aggregate information. Although it was interesting to know that 27% of a person's activity was browsing the internet, on its own this information was not rich enough to guide the development of new or better interface designs and applications. Over the years we have refined our logging methods to gather increasingly detailed information about user behavior. Our current logger, PersonalVibe[1], records millisecond-accurate start and end time stamps for each window activation on a user's desktop. For each active window it records title changes, executable information, an associated document if appropriate and a thumbnail. It also stores the start and stop times for input device use. As our

logger has evolved, so have our visualizations that we use to understand users' behaviors.

Two prior systems have been particularly inspirational in this work. TimeScape by Rekimoto [4] was a system that allowed the user to "unwind" the file system as a means to recall past activity. The system automatically aged documents from the desktop and when they were needed, the user could use the interface to "travel" back in time and view them. However this work only concerned itself with file creation and deletion and only used file creation time as a means of organizing files - no other attributes were available. MyLifeBits by Gemmell et al. [2] was a system that captured nearly all documents the user interacted with. MyLifeBits supported multiple tags and thus could handle both user authored tags and multiple time-based tags. However they did not pay attention to the user activity as a component. While the system would let users add multiple tags to any specific document, it did not provide support for discovering interesting temporal attributes of a file, such as when the document was used or what other documents were open at the same time.

Attempt 1

Our first attempt in representing users' desktop activities was a visualization based on a horizontal time axis showing input device activity and selected active windows (Figure 1). The colored bars represent the keyboard and mouse activity for an interval (the axis was user scalable). Below the axis is a thumbnail view of the most used application for that time period. The dark gray callout box displayed a list of the documents that were active during a particular interval. Documents could be selected from the list to re-open them.

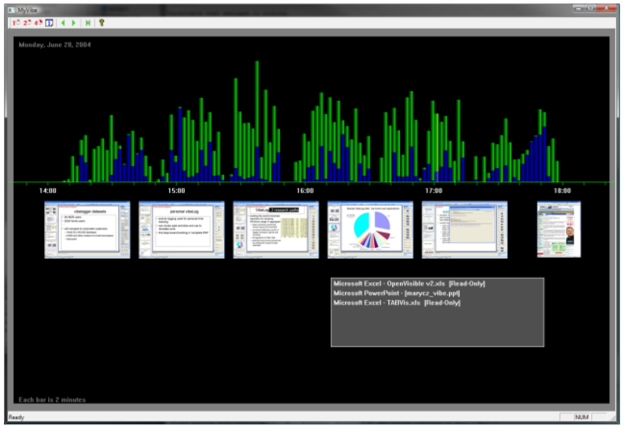

Figure 1: The first attempt to display window activation data.

There were two main problems with this visualization. First, the keyboard and mouse are not accurate proxies for the user's mental model of when he or she was busy. Thus the bars did not help people find the information they were looking for. Second, by only showing the most used window in the thumbnail area, many useful windows were hidden from the user. In fact, there was no way to discover those windows unless the user carefully moved the mouse and viewed the contents of the callout list at each interval.

Attempt 2

In Attempt 2 we removed the mouse and keyboard component and stacked the thumbnails in an attempt to address issues raised by Attempt 1. As can be seen in Figure 2, this resulted in a column of thumbnails that quickly exceeded the window height. To address this problem we overlapped the thumbnails and added a callout (shown with the dark black border) so that users could see the full thumbnail by moving the mouse

Figure 2 Second attempt showing only web pages loads.

over it. This approach was an improvement over the first but overall it presented more information than could be reasonably understood. One reason was that average activation time for any desktop window was a mere 20.9 seconds [3]. This rapid switching resulted in so many items that the display became visually overloaded. In fact, after talking to users and examining the data, it became clear that somehow the priority of each window needed to be conveyed better and that unimportant windows needed to be excluded from the visualization.

Attempt 3 – StatusWriter

Based on lessons learned from the previous two attempts, we wrote StatusWriter [1], shown in Figure 3. In this design each column represents a single day. The thumbnail representing each window is sorted by the sum of the total daily activity and placed on a tile that is drawn so that the more daily activity it has, the more it is "raised" above the background. To make up for the lack of a minute by minute timeline, a playback feature was also included. When a day is selected, the user can have a miniature desktop appear and replay the layout of all the windows for the day at various playback speeds. By using total active time per day as a proxy for the importance of a window we were able to get a fairly useful high-level, filtered representation of a day's activities. Most of the time, users were satisfied that the windows they wanted to view were present in the interface.

We interviewed several users about how they might use the interface to find documents and how they would use it to recall past work while writing a status report. Although StatusWriter proved fairly popular, there were

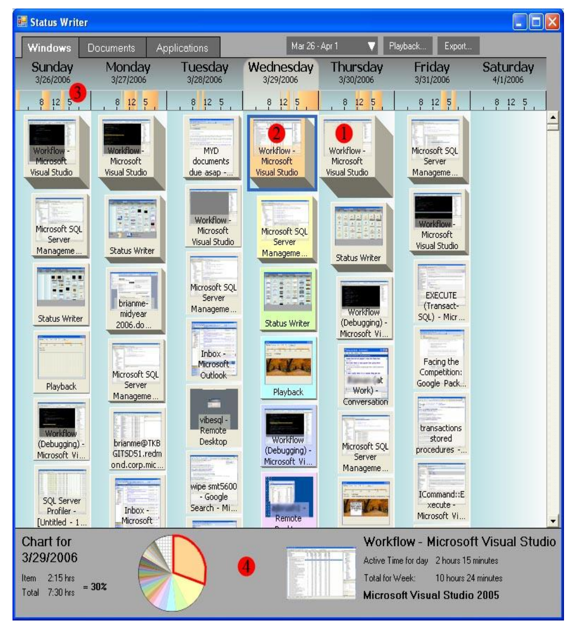

Figure 3: StatusWriter showing a week of computer activity.

again several problems that became clear during use. First, collapsing to a single day was not always the most appropriate level of fidelity for users. Often times, projects that took only a few hours were lost from view. Second, the absence of a timeline in the primary view proved to be problematic. Although seeing the summary for the day was helpful, users felt that they also needed to see the data on a timeline in order to find the documents they needed or be able to recall the day's events. Finally, users felt that they needed their calendar events represented on the same view as their computer activity.

Next Steps

At this point, we have a fairly clear understanding of the type of data we have available and which parts the users would like to see visualized. The problem is how to present the number of events that occur onto a timeline such that the user can view both the necessary details of a day, as well as an abstraction that would allow them to understand their computer behavior meaningful ways. Placing calendar events, which typically last an hour, onto a timeline with several thousand window switching events, which typically last 20 seconds, has so far been a substantial challenge. We've tried simply optimizing the zoom to make it possible for a user to quickly go from one level of detail to another but that still doesn't easily allow the user to have the overall daily activity context while they are at the zoom level that allows viewing individual window activity.

Currently we are prototyping interfaces that attempt to show multiple levels of detail by collapsing the high frequency window data. We are removing very short intervals and collapsing patterns into higher level objects. Although this shows promise we still can't reliably bring the thousands of window events into line with the human perception of "spending the morning working on a paper".

References

[1] Brush, A. B., Meyers, B. R., Tan, D. S., and Czerwinski, M. Understanding memory triggers for task tracking. In *Proc. CHI 2007*, ACM Press (2007), 947-950.

[2] Gemmell, J., Bell, G., Lueder, R., Drucker, S., and Wong, C. MyLifeBits: fulfilling the Memex vision. *In Proc. Multimedia*. ACM Press (2002), 235-238.

[3] Hutchings, D. R., Smith, G., Meyers, B., Czerwinski, M., and Robertson, G. Display space usage and window management operation comparisons between single monitor and multiple monitor users. In *Proc. AVI 2004,* ACM Press (2004), 32-39.

[4] Rekimoto, J. Time-machine computing: a time-centric approach for the information environment. In *Proc. UIST*, ACM Press, (1999), 45-54.

From Time to Time: User Modeling Through Temporal Data[1]

Emile Morse
National Institute of Standards and Technology
100 Bureau Drive
Gaithersburg, MD 20899 USA
emile.morse@nist.gov

Theresa O'Connell
National Institute of Standards and Technology
100 Bureau Drive
Gaithersburg, MD 20899 USA
theresa.oconnell@nist.gov

CHI 20089, April 4 – April 9, 2009, Boston, USA
ACM 978-1-60558-012-8/08/04.

Abstract
In this position paper we examine building user models and task models using temporal data. We discuss a large set of user interaction data collected at the National Institute of Standards and Technology using Glass Box software. We propose discussion topics on using temporal data from such data sets to build user models.

Keywords
Information visualization, temporal data, usability engineering, user interaction

ACM Classification Keywords
E.m [Data]: MISCELLANEOUS H.5.2 [Information interfaces and presentation]: User Interfaces—*Evaluation/methodology*. H.1.2 [Models and Principles]: User/Machine Systems—Human factors.

Introduction
Although people are the sum of what we know, at different points in time, the sum of what we know varies and so, any one of us can be many different people in terms of our knowledge.

User actions tell the story of who the user is at different points in time and what task the user is performing at a point in time. These actions are captured in data as entities. These actions, captured as data points, can inform user modeling and task analysis. Using data, in particular, temporal data, in user modeling and task analysis is the topic on which we would like to open a dialog.

Looking at temporal interaction data to build user and task models

Data are merely representations of information. As such, data can reflect relationships among information points. Bodies of information can change, e.g. they expand and diminish as new information is added or old information is discredited and discarded. As the information changes, the data change. In a data set where the data constantly changes, relationships among the data points are also ever changing.

In human-computer interaction (HCI) there is always a temporal aspect to the data that relates it to the task. There is a history to the data and to the relationships among the data points. At different points in time, the data shows us a different user profile. The user is the sum of all these profiles, slightly different at different points of time.

Task analysis consists of understanding the sequence of activities that comprise discrete tasks as well as the sequence of tasks that comprise coordinated or multi-step tasks. Understanding the temporal relationships among data points contributes to user modeling by informing task analysis. While human observation is a critical tool, it is our experience that in order to understand the complete story of user interactions, it is integral to include programmatic logging that captures temporal interaction data.

It is the role of usability engineers (UE) to understand who a user is in the context of the user's knowledge at different points in time while that person interacts with software. To accomplish this, the UE must examine the temporal aspects of interaction data. The complexity of such a study is well documented [e.g., 3, 4]. A temporal sequence of data points tells only one part of the story. The UE also needs to be able to jump among temporal points in order to gain the understanding of user activities necessary for user models and task models. In addition, the UE must be able to harness synchronicity of data points, particularly when users are engaged in collaborative activities. Even temporal intervals in the data are part of the story because they can inform identification and analysis of the points at which analysts examine evidence or simply stop interacting with the software so that they can focus on other cognitive aspects of analysis.

Because over time, interest and importance change, we require time-stamped user data. There are many tools to help us collect it, but they only bring part of the story to light. There are many types of relationships among the concepts that the data represents, such as geographic and historical relationships. For example, Oculus GeoTime shows both time and space very simply. Applying the GeoTime metaphor to other relationships among data points, can imply other relationships, e.g., a social network [1]. We have come to realize that while such tools are valuable, they are empowered by the data they use; their effectiveness is driven by the completeness, accuracy and comprehensibility of the data. And this data is augmented by UEs observing people in the context of using the software under study.

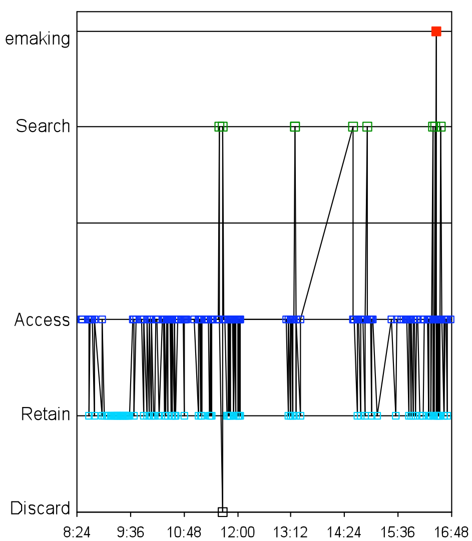

ure 1: Sample of 554 points of interaction

Glass Box data

At the National Institute of Standards and Technology (NIST) we have collected user data in large quantities over a variety of user activities. In particular, we have collected data from information analysts interacting with visualization tools. Originally, we used Glass Box software [5] developed at PNNL (Pacific Northwest National Labs). Over a period of 4 years, we conducted formative and summative usability evaluations using systems instrumented with the Glass Box. We collected hundreds of hours of user interaction data that included keystrokes, mouse moves and clicks, application invocations, document accesses both from the local file store and from Internet sources, copy/paste events, and other events as well as screen capture. All the events were time-stamped to the level of milliseconds. Reports of the results were shared with teams developing research-grade software and some results have been presented [5].

Based on our experience with the Glass Box, our team was able to determine a set of critical events that needed to be captured to adequately represent the analytic process used by intelligence analysts. And based on our interest in moving from a thick-client to a Web client, we worked with a group of researchers to develop a specification of a new method for logging time-stamped user interaction data. The critical events include:

- Search – user issues a query usually to an Internet search engine

- Access – user opens a document either from a local file-system or using a browser

- Retain – user copies text from one document and inserts it into another location

- Sensemaking – user constructs hypotheses and marshals evidence that supports or contradicts that hypothesis

- Discard – user removes an information object (e.g. query, piece of evidence, hypothesis, or document)

Figure 1 shows data collected while a user was working on an intelligence analysis task. The user's job was to find information relevant to a topic and to produce a report describing his hypotheses, the evidence for and against each of those hypotheses and his conclusions. The figure shows the analyst's work on the third day of a four day assignment. The gap in the middle maps to the lunch break. The single sense-making event toward the end of the day was the first such event of this type. The basic pattern reveals lots of reading (Access) alternating with cutting and pasting information (Retain). We employ a variety of methods to detect interaction patterns in the data in an iterative process of exploratory sequential data analysis [3]. Detecting and explaining gaps (e.g. lunch), finding outlier events (e.g. the sense-making data point), and noting repetitive patterns (e.g. Access-Retain cycling) are some of the important factors in analyzing our users' behavior.

Having time-stamped capture gives UEs a broader view of the users' knowledge state, i.e. a broader view of who the user is at any point in time. This is an example of letting systems do what they're good at – getting the minute level right – while letting people do what they excel at – getting the gestalt right.

To accomplish user profiling through mining temporal data is a multidisciplinary task, e.g., it is the work of psychologists, usability engineers computer scientists and social scientists, to discern user models and user tasks. It is our hope that people from all the professions involved in understanding how humans interact with computers will join in discussing the use of temporal data in building user models.

Citations

[1] Chien, L., Tat, A., Proulx, P., Khamisa, A., Wright, W. Grand challenge award 2008: Support for diverse analytic techniques - nSpace2 and GeoTime visual analytics. Proceedings of IEEE VAST Symposium, (pp. 199-200). Piscataway, NJ: IEEE. October 21-October 23, 2008.

[2] Cowley, P. Nowell, L. Scholtz, J. Glass Box: An Instrumented Infrastructure for Supporting Human Interaction with Information. HICSS '05. Proceedings of the 38th Annual Hawaii International Conference on System Sciences, 2005. 296c - 296c

[3] Fisher, C. and Sanderson, P. 1996. Exploratory sequential data analysis: exploring continuous observational data. *interactions* 3, 2 (Mar. 1996), 25-34.

[4] Hilbert, D. and Redmiles, D. 2000. Extracting usability information from user interface events. *ACM Comput. Surv.* 32, 4 (Dec. 2000), 384-421.

[5] Scholtz, J., Morse, E., and Hewett, T. 2004. In Depth Observational Studies of Professional Intelligence Analysts. Human Performance, Situation Awareness and Automation Technology Conference.

[1] Certain commercial equipment, instruments, materials, services or companies are identified in this paper. This in no way implies endorsement or recommendation by NIST.

Temporal data representation on mobile devices for in-field law enforcement

Avin Pattath
Purdue University Regional Visualization and Analytics Center, Purdue University
465 Northwestern Ave.
West Lafayette, IN 47906 USA
apattath@purdue.edu

William Pike
Pacific Northwest National Laboratory,
902 Battelle Boulevard,
Richland, WA 99352 USA
william.pike@pnl.gov

David Ebert
Purdue University Regional Visualization and Analytics Center, Purdue University
465 Northwestern Ave.
West Lafayette, IN 47906 USA
ebertd@purdue.edu

CHI 2009, April 4 – 9, 2009, Boston, MA, USA
ACM 978-1-60558-246-7/09/04.

Abstract

Law enforcement operations produce and consume temporal data. Presenting such data on portable devices to enable in-field decision making is challenging. In this paper, we present a classification of law enforcement operations based on their task requirements and response times followed by a discussion of issues in temporal data representation for each task type.

Keywords

Temporal data visualization, mobile visual analytics, mobile emergency response, in-field response and investigation

ACM Classification Keywords

H.5.2 [INFORMATION INTERFACES AND PRESENTATION (e.g., HCI)]: User Interfaces---User-centered design, Miscellaneous; J.7 [COMPUTERS IN OTHER SYSTEMS]: Command and control, Real time; I.3.6 [COMPUTER GRAPHICS]: Methodology and Techniques---Interaction Techniques.

Introduction

Temporal data representation and interaction on mobile devices presents unique challenges arising due to the limited form-factor and situational context of usage of these devices. Mobile devices have been traditionally used for personal data management to represent time-related data such as calendars and task schedules. However, with their ever-improving computational power, they are increasingly being used for critical response and investigation operations in the field [1,4,5]. These operations produce and consume temporal data from disparate sources and in various forms. While the backend representation of time-related data is similar to that of desktop based systems, client side data representation should consider the form-factor and situational constraints of the mobile device. In this paper, we specifically focus on the issues of representing time-based data encountered in in-field law enforcement operations.

In-field law enforcement operations

Tracking (people, objects), patrolling, evidence and data collection, situation monitoring and investigative analysis are some in-field operations performed by law enforcement officers that can benefit from using a mobile device for making rapid decisions. Due to varying requirements and response times for each task, combined with the device constraints, it is necessary to adapt the representation and notification of temporal data accordingly.

We classify the major task types in this domain and later discuss issues in temporal data representation for each task type.

Classification of law enforcement operations

We classified various law enforcement operations into three major categories:

- In-field response (real-time and historical data)
- In-field investigative analysis (real-time data)
- Post-event investigative analysis (historical data)

In-field response

In-field response refers to operations that involve immediate action. Such operations could result from real-time data obtained during patrolling, evidence gathering and tracking operations or from historical intelligence data. Real-time data obtained from sensors such as GPS (Geographical Positioning System) and cameras, are updated in intervals of a few seconds. For response operations, it is important to provide accurate information at the appropriate frequency using effective representation and notification cues. Real-time data can be obtained by either continuously or periodically polling for data from the sensors. However, considering the limited battery resources on a mobile device, we have to poll for data periodically while sacrificing the recency of data. Entin et al. [3] observed that users in fact perceive movements better when locations of their team members are updated periodically rather than continuously. Recent updates of information should be represented in the context of their past, so that in-field responders can track changes over time. However, considering the limited memory on mobile devices, we need to strike a balance between the amount of history stored (and displayed) and the available memory on the device. Typically, such data is displayed using ghosting techniques as used in [4] with the faded trail region indicating data from the recent past. In-field responders also may often have to turn their attention away from the device towards other activities. Given their limited cognitive resources while performing an operation, it is necessary to "push" important updates to the client in the form of audio or tactile notification cues [2] when they arrive. Such "push" based mechanism is also effective for presenting temporal or spatial context-based notifications generated by background monitoring processes from historical intelligence data.

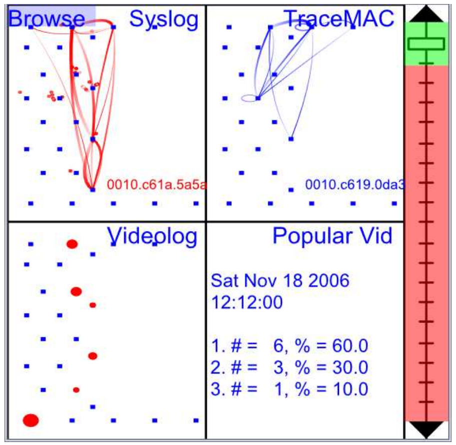

figure 1. A screenshot from the NetworkVis system [5] showing aggregated overviews of temporal data with tools for temporal analysis.

In-field investigative analysis

Investigative analysis in the field involves determining emerging trends and patterns during an event and taking appropriate action. Typically, such actions are performed by event monitors and supervisor patrols (who are responsible for the placement of their team members). Real time data is used to detect trends or anomalous patterns as the events unfold. However, such trends can be detected only after obtaining a significant amount of data. Hence the frequency of updates needs to be slower than that for in-field response. Moreover, the response time is also greater in this case. The time scale and data representation should reflect these requirements. For example, fig. 1 shows a screenshot from a system (NetworkVis) that we developed [5] to monitor and troubleshoot network activities at our university's football stadium during home game days. The system was setup to pull real-time data at regular intervals from the network log database (the green color on the timeline shows the time until when data is available). The time scale was also set to the same interval and we represented all time dependent data at two levels. The overview level shown in fig. 1 displays only the most important information aggregated from various sources (such as access points, video and game servers), to compensate for the small screen size. The second level provides detailed information upon demand and according to the analysis needs. Interacting with the timeline allows the analyst to browse through time and detect unusual patterns in the network.

Post-event investigative analysis

Post-event analysis, such as incident and crime scene analysis, is usually performed by specialized analysts in the field after an event occurs. While this type of analysis does not usually involve real-time data, it usually accesses data ranging across a broader time scale. Therefore it is important to provide temporal analysis tools at multiple granularities to discover patterns at different time scales. Moreover, it is important that the user be able to intuitively interact with and explore the data. Figure 2 shows a screenshot from a mobile system we are currently working on, which provides interactive tools to analyze crime incident reports in a small geographic region. It uses a focus + context exploration lens that also doubles up as a spatial filter for geo-tagged data. The current version just supports analysis at a single temporal granularity. We calculate simple statistics (bar graphs) on the fly to indicate number of crimes in each category. Moreover, it is also useful to display crime graphs varying with time at various granularities to provide insight into their trend variation.

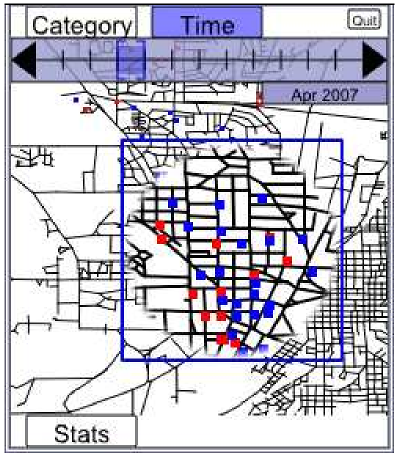

figure 2: A screenshot from our ongoing work for in-field crime incident analysis showing temporal and geospatial data

An important issue we wish to discuss in this workshop is the selection of appropriate time-scale granularities for various tasks that decide the level of corresponding data aggregation. Since this is highly task-specific, as a first step, we classified common tasks in our domain and described their temporal data characteristics. Screen space and cognitive resources are limited while using a mobile device. Therefore, data aggregated at the appropriate level should not only convey meaningful information, but also make it easily perceivable.

Acknowledgements
We wish to thank the U.S. Department of Homeland Security's Regional Visualization and Analytics Center (RVAC) program for funding our work.

References

[1] Baddeley, B., Younkin, K., Riensche, R., Best, D., Pike, W., and May, R. From desktop to field: Deploying visual incident analysis for law enforcement. In *IEEE Conference on Technologies for Homeland Security 2008,* 1-8.

[2] Campbell, C., and Tarasewich, P. Creating notification cues for mobile response teams. In *Proc. Of 2004 DSI Annual Meeting.*

[3] Entin, E. B. Representing and Visualizing a Dynamically Changing Tactical Situation. In *Proc. IEA 2000/HFES 2000 Congress,* 427-430.

[4] Kim, S., Jang, Y., Mellema, A., Ebert, D., Collins, T. Visual Analytics on Mobile Devices for Emergency Response. In *IEEE Symposium on Visual Analytics Science and Technology (VAST 2008)*, 35-42.

[5] Pattath, A., Bue, B., Jang, Y., Ebert, D., Zhong, X., Ault, A., and Coyle, E. Interactive Visualization and Analysis of Network and Sensor Data on a Mobile Device. In *IEEE Symposium on Visual Analytics Science and Technology 2006,* 83-90.

Interactive trace visualization in Trace-based systems

Julien Laflaquière
Université de Lyon, CNRS
Université Lyon 1, LIRIS,
UMR5205, F-69622, France
Julien.laflaquière@liris.cnrs.fr

Yannick Prié
Université de Lyon, CNRS
Université Lyon 1, LIRIS,
UMR5205, F-69622, France
yannick.prie@liris.cnrs.fr

CHI 2009, April 4 – 9, 2009, Boston, MA, USA
ACM 978-1-60558-246-7/09/04.

Abstract
We first introduce the notions of modeled-traces (M-traces) and trace-based systems. We then present activity analysis and activity support as the two main uses of M-traces, and we discuss interactive visualization challenges related to active reading of traces.

Keywords
Activity analysis, activity support, modeled-traces, trace-based systems, trace visualization

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous. See [3] for help using the ACM Classification system.

Introduction
Numerous research works have investigated the general issue of *use traces*. Such temporal data, sometimes called interaction history or user data, are collected during the interaction between a user and a digital environment. The collected traces are most of the time composed of *log files* related to the system's functionalities. Those files being usually too complex for human direct interpretation, statistical treatments are usually needed for their rendering, aiming at converting

quantitative data into qualitative visual insights that characterize the usage, the user or the activity. In that case, visualization techniques are applied directly on raw data so as to make remarkable patterns emerge [2].

Trace-based systems

The main goal of our approach is to be able to manage and present digital traces that are qualitative representation of observed activities. We argue that a means to reach such a goal is to consider an *explicit* and *activity-oriented* modeling of traces. 'Explicit' here opposes to the implicit modeling of most of the log files, while 'activity-oriented' opposes to their function-oriented (or design-oriented) modeling.

In this paper we deal with the visualization of M-traces that have been transformed so as to be used in particular situations.

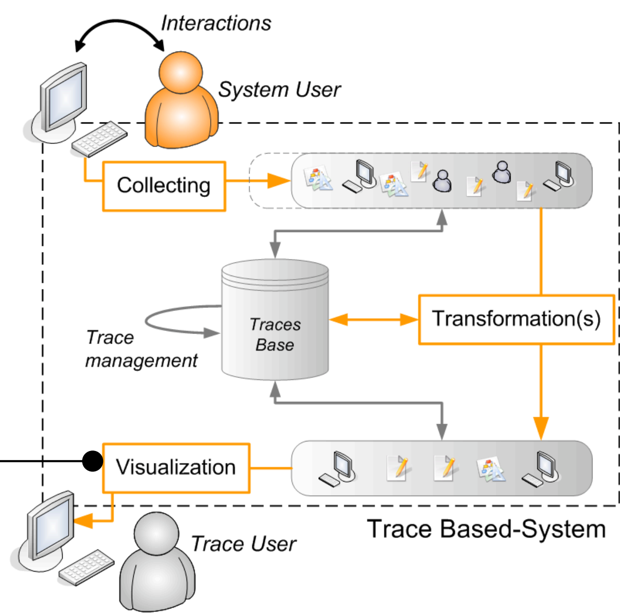

Figure 1: Trace-based system global architecture

Our approach at the LIRIS-SILEX team is conceptualized in the *Trace-Based System* framework [1]. At the core of this framework is the notion of *modeled-trace* (or M-trace) defined as the association of a *collection of temporally situated observed elements structured by relations* with an *explicit model of these observed elements and relations*. Such explicit trace models can be considered as ontologies of M-traces that describe both their constitutive elements (entities or events) and the complex relationships that can be expressed between them.

Though it can easily represent simple event log files, the notion of M-trace can also tackle complex application use histories. Moreover, being explicitly modeled with knowledge representation techniques, traces can be managed in generic and flexible ways. They are considered as first class objects that can be manipulated independently of the particular environment that was involved in the activity they are traces of.

A trace-based system is then a system dedicated to the management of modeled-traces, from their collecting to their use. It offers services for *trace transformations* such as fusion, filtering and rewriting (figure 1). Such transformations can help abstracting traces from low-level *a priori*, log-like modeling to higher-level *a posteriori*, activity-oriented modeling. Abstracting can be important for reaching an adequate level of interpretability, especially in the case of visualization.

Interactive visualization of M-traces

In this article we focus on two main situations for M-trace exploitation: *activity analysis* and *activity support*.

Activity analysis
In a situation also known as *Exploratory Sequential Data Analysis* [3], an analyst who was not involved in the traced activity transforms an M-trace so as to characterize the investigated activity. It is the case in the *Abstract* approach [4] where a TBS was implemented for facilitating the analysis of behavioral data in a vehicle-driving situation. The analysis is carried off-line, after the action, thanks to a visual interface (figure 2). On the contrary, in the *eLycée* project [5], the M-trace is displayed during the activity. Students work in a virtual classroom with a tutor while their activity is traced. Such traces can be exploited after the class session but also by the tutor directly *during* the session, so as to manage the group activity online. The challenge here is to provide a M-trace that describes not just what is going on but *how* students are or were behaving, during or after the time of their activity.

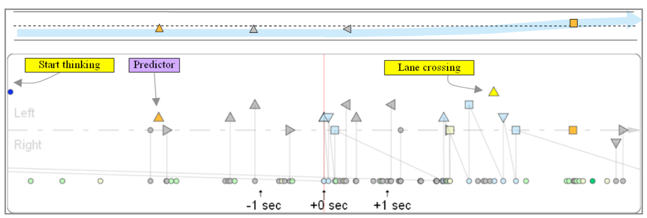

Figure 2: Three abstraction levels of driving-related M-traces in the *Abstract* system [4].

Activity support
M-traces are here directly used *within* the activity, the user of the M-traces being the user of the environment they are collected from. In that case, traces are used for facilitating the activity by visually representing how it takes or took place be it in real time or afterwards. We are here interested in complex activities the unfolding of which cannot be fully planned in advance, such as video active reading with Advene [6] or the production of multimedia learning material with a dedicated software (*Emulsion*).

Figure 3: Using MIT SIMILE timeline for *Emulsion*.

In such cases, we try to facilitate the activity by providing the user with a rich representation of her own activity 1/ so that she can get a sense of her current situation; 2/ for contextualizing the elements she manipulates; or 3/ for reusing past significant episodes so as to produce new activity phases.

Some interaction challenges

Though an M-trace is modeled specifically for a dedicated type of activity so as to ‘make sense’ with respect to it, no *a priori* pre-determined visualization can pretend to match all its interpretation needs. This is true both for activity analysts and users observing their own activity, and it appears to us that the solution consists in the interactive visualization of traces. We claim that the interpretation itself must be constructed along such interactive visualization, by *trace active reading*, which poses several challenges.

Model-based presentation of traces
M-traces are traces whose elements and relations are described abstractly in knowledge models. Such models can be used as inputs for model-based trace visualization interfaces that give user rich access both to abstract (model) and concrete (actual trace(s)) knowledge structures.

Time-based and structural presentation of elements
Most interfaces for traces visualization are currently based on the temporal dimension, emphasizing mainly the temporal relationships between observed elements. Horizontal or vertical timelines present punctual or durative observed elements, often neglecting the structural relationships between them. A challenge here is to design interactive interfaces that smoothly articulate and present at the same time temporal *and* structural relationships between elements.

Panoptic organization vs. rich elements representation
Any timeline designed for trace visualization needs to cope with two contradictory needs. First the need for a synoptic vision of a slice of activity, hence for a light visualization of elements within their time range. Second the need for a rich contextualization of elements that are isolated within the trace, hence for a rich visualization of elements. An important challenge here concerns interfaces suited to adapting temporal elements visualization to temporal range visualization, while providing users with means to control and tailor such adaptation.

Interfaces for trace filtering and rewriting
Interactive visualization of traces means providing the user with the full control of the representation of her activity. Here TBS generic trace transformation services can help her construct traces that match her representation needs. *Filtering* offers choice of which elements of the trace model (actions, events, resources…) are to be presented. *Rewriting* provides the user with the ability to create her own trace model, by determining *new* types of observed elements, and how they are to be constructed from existing ones. Numerous challenges are related to so defined trace active reading: How to interact with a trace or a trace fragment and specify its filtering or rewriting? How to manage one's multiple traces and transformations? How to elicit new relationships within a trace?

References

[1] Laflaquière, J., Settouti, L.-S., Prié Y. and Mille, A. A trace-based System Framework for Experience Management and Engineering. *2nd Int Workshop on Experience Management and Engineering* in conjunction with KES (2006), 8 pp.

[2] Chen, J., Sun, L., Zaiane, O. R. and Goebel, R. Visualizing and discovering web navigational patterns. *7th Int. Workshop on the Web and Databases, ACM SIGMOD/PODS 2004 (2004),* 13–18.

[3] Fisher, C. and Sanderson, P. Exploratory Sequential Data Analysis: Traditions, Techniques and Tools. *ACM SIGCHI Bulletin*, 25, 1 (1993), 34-40.

[4] Georgeon,O., Mille, A. and Bellet, T. Analyzing behavioral data for refining cognitive models of operator. *Philosophies and Methodologies for Knowledge Discovery*, *7th Int. Workshop on Database and Expert Systems Applications* (2005), 588-592.

[5] Cram, D., Jouvin, D. and Mille, A. Visualizing Interaction Traces to improve Reflexivity in Synchronous Collaborative e-Learning Activities, *6th European Conference on e-Learning* (2007), 147-158.

[6] Aubert, O., Champin, P.-A., Prié, Y. and Richard, B. Canonical processes in active reading and hypervideo production. *Multimedia Syst.*, 14, 6 (2008), 427–433.

Interacting with annotated videos

Yannick Prié
Université de Lyon, CNRS
Université Lyon 1, LIRIS,
UMR5205, F-69622, France
yannick.prie@liris.cnrs.fr

CHI 2009, April 4 – 9, 2009, Boston, MA, USA
ACM 978-1-60558-246-7/09/04.

Abstract
We present Advene, a project and a prototype dedicated to video active reading, i.e. inscribing annotations on a video document and taking advantage of these annotations for navigating it. Both the video and its annotations are temporal data that need to be presented and interacted with. We present some illustrative interfaces, and discuss some of the challenges related to such a complex activity.

Keywords
Audiovisual annotation, active reading, document time, stream time, video interaction

ACM Classification Keywords
H.5.1 [Multimedia Information Systems] : Video; H.5.2 [User Interfaces] : Interaction styles, Theory and methods; H.5.4 [Hypertext/Hypermedia] : User issues

Introduction: video active reading
Video documents take more and more importance within the digital world. Of course, their most widespread use is related to their basic consumption from beginning to end thanks to basic video players, the main useful commands of which are *play* and *pause*. Nevertheless, some video-related activities go beyond such basic consumption and need *deep manipulation* of videos. Amongst those activities are

montage (creation of a video from several video and audio tracks) and active reading.

Active reading is a so-called “knowledge-work” activity that consists in taking notes on documents while reading them, so as to be able to produce new documents. Such notes are called annotations; they usually feature an anchor (link to a fragment of the document) and content. In audiovisual active reading, annotations anchors usually are temporal fragments of the stream (defined with two time codes), while their content is most often textual. New documents resulting from active reading sessions fail into the general category of hypervideos. Video active readers include ergonomists, linguistic or gesture analysts that annotate video captures of activities for analyze; film enthusiasts or critics who annotate movies so as to publish and share their analyses as hypervideos.

The Advene project and prototype

The Advene project (Annotate Digital Video, Exchange on the Net, http://advene.org) aims since 2002 at studying the emergence of hypervideos – as new hypermedia forms created upon audiovisual documents – and at designing systems that allow creating, building and sharing hypervideos [2]. It features a data model [3] where annotations, relations between annotations, description schemes, and views definitions related to one video are grouped into a container named “package”. Views can be either static (web-based, e.g. a table of contents in a browser) or dynamic (player-based, e.g. subtitling of the movies during playing). The set of views related to a video constitutes a hypervideo that can be shared by sharing the package (annotation structure, view definitions) it is generated from.

The Advene prototype [1] is a free, python-based, multi-platform, open source application that includes a video player, and offers active reading facilities, implementing the Advene data model. It allows the user to annotate video documents and to define both static and dynamic views so as to express her personal analysis that can be further shared and refined. It also features several GUIs that are used for annotating the video stream and managing its annotations.

Two kinds of temporal data

There are two kinds of temporal data that need to be presented and interacted with during video active reading. First, the audiovisual document itself is natively temporal, which means that it mandatory needs a player to be played. Two temporalities are related to such documents: the *stream temporality* corresponds to the stream being played at each instant at a certain rate, it is the machine temporality; the *document temporality* comes from the duration of the document and signifies that is has a beginning, an end, and a certain length. Second, the annotation structure is composed of interrelated annotations. Each annotation is anchored in the temporal stream, inheriting its temporality from its anchor fragment. A set of annotations constitutes temporal data over the video document.

Interacting with the video to annotate it and with annotations to access the video

We have designed several GUIs so as to be able to carry active reading activities. Some of them are presented on figure 1. The most prominent interfaces are a video player to (dis)play audiovisual information, and a classical timeline that presents document time over annotations categories. Trough the representation of the document temporality, it offers navigation within the document; presenting the annotation structure it allows its creation and management through various interaction modes.

The treeview is a classical hierarchical non-temporal presentation of the information contained in the opened package.

The note-taking view allows the user to take note over the video stream by inserting time stamps inside her text, making it an interactive access to the video. Time-stamped text can further be converted to annotations.

The active bookmark view allows taking snapshots of the video stream and further describing them. It is then possible to combine snapshots (time stamps) into annotations (temporal fragments).

The note taking view adds temporal markers (here represented as images) to a sequence of characters, thus defining annotations fragment and content.

The Active bookmarks view offers direct manipulation of bookmarks and annotations as pairs of bookmarks.

The annotation structure is presented in the timeline view. Abstract annotation types give vertical organization. Relations (here *dialog > image*) permit to express more semantics that just plain annotations. The spatialization of the annotations temporalities allows simple manipulation of complex data (e.g. setting the end time of an annotation as the end time of another).

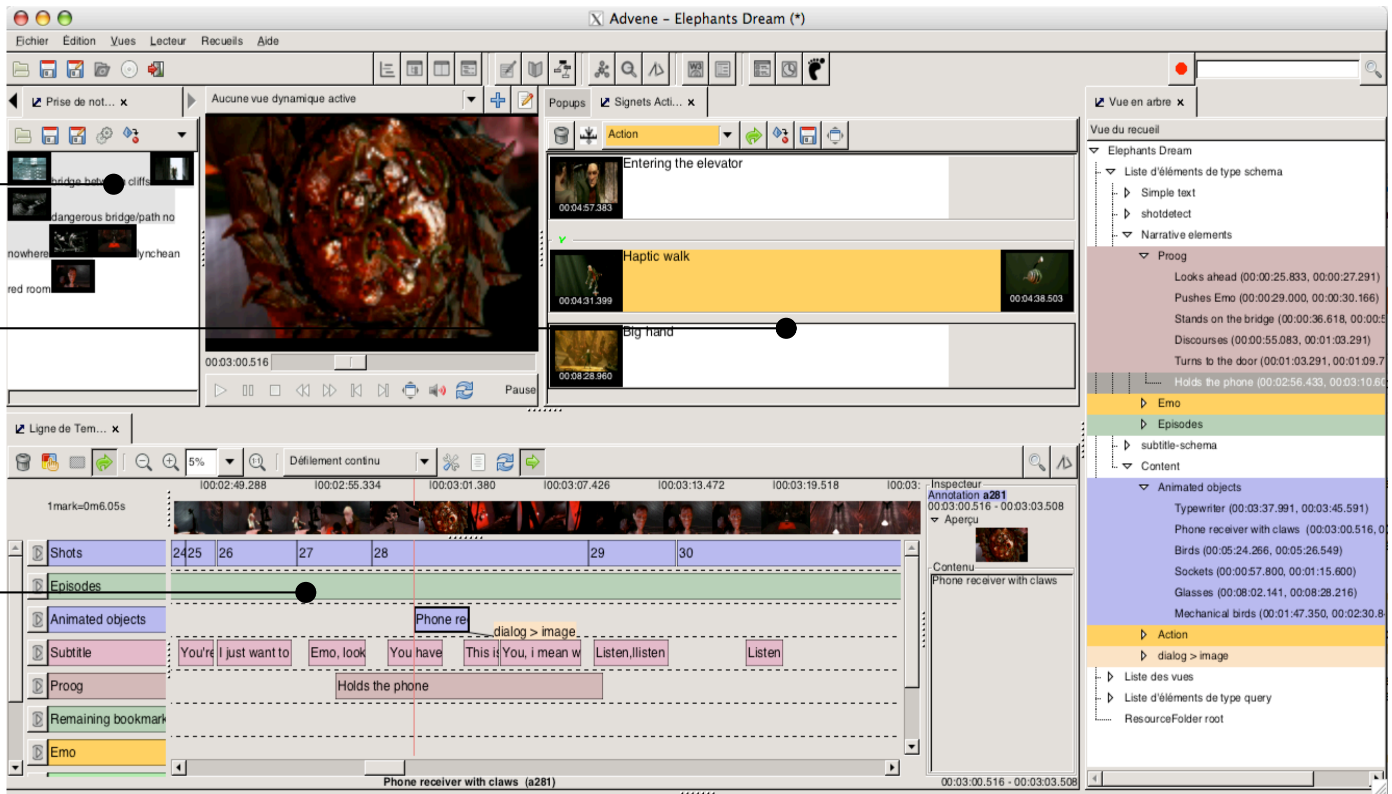

Figure 1. The Advene main interface with several *ad hoc* views: timeline, treeview, note-taking, active bookmarks.

Interaction challenges

The interfaces we have rapidly presented illustrate several challenges that are related to interacting with video annotations.

The first challenge consists in being able to *interact with an audiovisual annotation structure* (structure = annotations + relations + description schemes), which is data that on the one side is inherently *temporal* (from its inscription into the stream through fragments) and on the other side has *atemporal* features, as a knowledge structure. A table of contents of a film, or the description of the diegetic relationships between characters in the movie can be considered both atemporal and temporal, whether they allow the playing of the stream, are presented in a temporal or a sequential manner, etc. An annotation can be (rightly or erroneously, depending on its author's intention) considered both as belonging to a static structure interpretable in the absence of the movie, and as being tightly connected to it, so that it cannot be interpreted without actually playing the related fragment.

As such playing takes time, a sub-challenge relates to *static representation of a temporal stream and its fragments* in the case of video. Video being visually pregnant media that can impress user, fragments can be presented as atemporal structures of excerpt images that can act for the user as mental indexes for her memory of the video stream. Beside classical use of one (begin) or two (begin + end) images for representing fragments, there is plenty of space for new ways of statically representing them, the document they belong to and their relations [5].

Our second general challenge concerns *the management of the temporality mix in the activity of active reading of a temporal stream*. Shortly presented, GUIs for video active reading have to take into account the *stream temporality* (and the associated fascination for the reader) and the *inscription temporality* (temporality of the annotation activity and of all its subtasks, e.g. create an annotation, adjust its begin or end time, type in its content, etc.). The major problem being that annotations are created as a consequence of the playing of a video, but that this very playing becomes annoying as soon as the annotation has to be inscribed, because the task of creating the annotation and the task of going on watching the video collapse. Important issues related to this second challenge relate to loss of attention and recovery, controlled task interruption and resumption [4].

References

[1] Aubert, O. and Prié, Y. Advene: active reading through hypervideo. *Proc. ACM* Hypertext'05 *(*2005), 235–244.

[2] Aubert, O., Champin, P.-A., Prié, Y. and Richard, B. Canonical processes in active reading and hypervideo production. *Multimedia Systems*, 14, 6 (2008), 427–433.

[3] Champin, P.A. and Prié, Y. Models for sustaining emergence of practices for hypervideo. *SADPI'2007* (2007), 59-68.

[4] Iqbal, S. T., and Horvitz, E. Disruption and recovery of computing tasks: field study, analysis, and directions. CHI '07 (2007), 677–686.

[5] Money, A.-G. and Agius, H. Video summarisation: A conceptual framework and survey of the state of the art. *Journal of Visual Communication and Image Representation*, 19, 2 (2008), 121-143.

Studying long-term, fragmented data sets

Stuart Reeves, Alistair Morrison and Matthew Chalmers
Computing Science,
University of Glasgow.
Glasgow G12 8QQ, UK
{stuartr, morrisaj,
matthew}@dcs.gla.ac.uk.

CHI 2009, April 4 – 9, 2009, Boston, MA, USA
ACM 978-1-60558-246-7/09/04.

Abstract
In this paper we suggest two challenges for the study of fragmented data sets generated from long-term studies. The first of these is the wide range of temporal perspectives from which a single data set may be inspected (from seconds to weeks). The second challenge involves the importance of considering user experience of time as a useful resource in analysis. Finally we briefly conclude with a call to consider new analytic tools that move beyond solely timeline-bound representations.

Keywords
Temporal data sets, qualitative and quantitative analysis tools, ubicomp systems

ACM Classification Keywords
H.5.2 User Interfaces: Evaluation/methodology.

Introduction
The evaluation of ubicomp systems often creates difficult demands both for the practicalities of data collection, and the evaluative practices of researchers during and after the period of use. The fragmented, socially embedded and typically mobile nature of ubicomp systems is coupled with the frequent desire of researchers to understand the ways in which

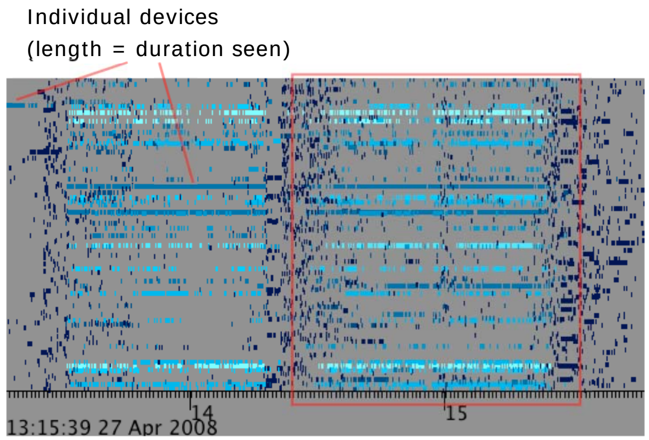

Figure 1: The density of several hours of Bluetooth devices around an individual, before, during and after a football match. The individuals that have been seen for the longest accumulative time are shown in a lighter colour. Within this data, game patterns, such as the two halves of the match (second half marked) and subsequent journey home are visible.

technology comes to be woven into everyday life and 'mundane' through adaptation, appropriation and contextualisation of system use. More often than not this means that periods of use under evaluation are as long-term as possible (although there are notable exceptions, particularly within ubicomp research that examines performance settings, e.g., [7, 1]), and, increasingly often, involves trials and deployments that are conducted 'in the wild' of everyday life rather than in controlled environments [6].

Temporal range

This concern with evaluation periods lasting days, weeks and sometimes months in natural settings, in combination with the disparate, transient and mobile forms of interaction that occur in many ubicomp systems, results in data sets fragmented both spatially and temporally [2]. Assuming the considerable problems of data collection are overcome, analysts must then reconstruct and piece together fragmented sets of diverse data types (e.g., audiovisual recordings, field notes, system transaction log files, sensor logs such as GPS or wifi scanning data, etc.). A subsequent issue that faces analysts is the potential range of temporal scales under which this data set can then be examined.

To illustrate this we can consider examples from our own and others' research projects. Recently, for example, we have been investigating the conduct of crowds at and around stadium-based sporting events. As part of this we have collected a body of video recordings documenting the forms of interaction that spectators engage in. Additionally, we have recorded Bluetooth device identity data using scanners carried by those attending these events, in an attempt to document the crowd's conduct at large. In studying this data we are interested both in the large-scale trends of crowd activity detected by the simple scanning

Figure 2: A small moment of interaction from a match-day video recording in which two sets of football supporters located in the same pub – fans of the Scottish and Norwegian national teams – 'break the ice', bridging between opposing crowd groupings via intimate gestures.

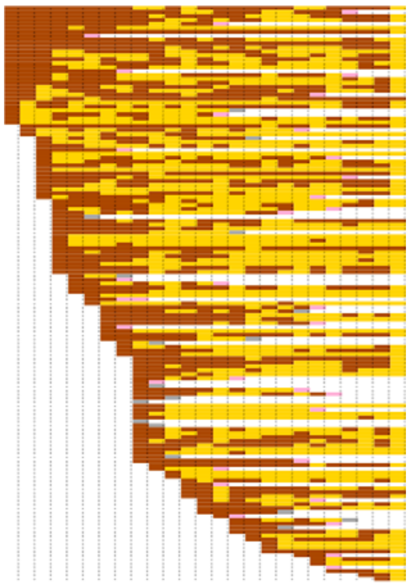

Figure 3: Engagement of players in Day of the Figurines over the course of 24 days (taken from [3]). Red bars indicate days players were 'engaged' in the game, yellow bars indicate that they were 'disengaged' and white indicating that they had exited the game.

technique (see Figure 1), and the smallest of interactional moments of crowd activity (e.g., see Figure 2). In spite of the differences, analysis of extreme ends of the temporal spectrum can often be mutually informing, providing context and richer understanding.

Analysis of other projects, such as Day of the Figurines (DoF), a narrative-driven interactive SMS-based game for mobile phones, reveals an even wider temporal scope, and further issues when examining temporal data. The game itself was intended to be played episodically, unfolding "in the background of players' daily lives" [5]. Once again, log data and video recordings collected from weeks of play by hundreds of players were examined both for long-term patterns of engagement (see Figure 3, covering 24 days) [3], as well as for fine-grained ethnographic analysis of orchestration practices [4].

Thus, we can see how an analysis of data collected from a ubicomp system trial may have stages that focus on very different levels or scales of temporal detail.

User experience

Interestingly, the DoF analysis also highlighted how players may experience the flow of time differently due to their varying levels of engagement in the game (e.g., playing every day versus playing once a week). Thus, when addressing temporal data sets, understandings of different layers of experienced "temporal trajectories" [5] becomes key. In DoF, the transactions between players and game orchestrators in collaboratively producing a narrative, and player engagement in the game (Figure 3) can be understood in terms of different temporal narrative trajectories. For analysts, therefore, not only is the sheer range of temporal scales potentially very wide for any relatively diverse data set, but also we can begin to see how the temporal experience of those being evaluated may vary greatly. The relevance of this may well need to be folded into the analytic methodology.

We might also consider this subjective temporal experience for shorter periods. For the football supporters in the first example, for instance, 90 minutes of the match experience will differ greatly in pace when compared to discussing the game with others subsequently. Pushing this notion further, we can consider more extreme analytic examples, such as the study of interactive systems that operate on fairground rides [1]. In this instance, as with many 'stressful' experiences, perception of time radically changes [12], and so interpreting sensor and video data of participants may well need to take this into account when evaluators try to make sense of these experiences.

Challenges

From these observations on the evaluation challenges posed by varied data sets and projects, we can make some more general comments on the broader challenges for the evaluation of fragmented, long-term data sets.

Firstly we looked at the potential for very wide temporal range within sets of data. The issue here then is what might be the best tools and methods for analysts concerned with differing ends of the temporal detail spectrum. We saw both aggregations of data over the course of hours, days and weeks, and how segments of (typically video) data from this same data set that may last a mere few seconds can also be of interest. For example, long term appropriation of a technology may hinge on many brief moments of interpersonal interaction.

In dealing with data at varying temporal scales, an analyst is likely to need a variety of different types of tools, each with varying strengths and weaknesses. For example, whereas it is possible to look at trial-wide overviews of log data to examine long-term trends of system use, it is harder to summarise the contents of video data in a useful way without sitting through the full corpus of collected footage. Driven by the recognition that varying types of recorded data form part of the same data set, one might ask whether it is of greater benefit to use a collection of tools, each specialised for a single type of data (often a single end of the temporal spectrum) or a system that attempts to merge such tools into one application, but the most likely way forward may be the linking of specialised tools to form one co-ordinated ensemble. A key design characteristic of information visualisation is to show fine-grained detail in the wider context of, ideally, the entire data set [8] Systems such as DRS [11] and our Replayer system [9] have taken this approach, aiming to show how fine-grained segments may elaborate upon the whole (and vice versa). We note, however, that a perennial design challenge with this approach is how to let users move smoothly between overviews of time and the most useful small periods within it. Drilling down, as such, is not a problem but knowing where to drill down is. For example, it would be technically straightforward to use the Bluetooth scanning data visualisation (Figure 1) as an index for video clips recorded by the individual conducting the scanning, but in practice the quantitative features that the visualisation shows may not be the most useful clues as to the qualitatively most significant short sequences of video.

A similar issue with wide-ranging fragmented temporal data sets is the often large differences in the rates of collection or sampling of quantitative data. For example, accelerometer and related sensors may need sampling rates of over 50Hz in order to get data accurate enough for some modelling approaches, but it is not immediately apparent how to visualise—or visualise well—that data alongside GPS readings being taken once every 10 seconds, or field notes taken every few minutes, especially if the patterns under investigation are happening at a very fine scale. Like the bridging between fine-grained video clips and coarse-grained Bluetooth patterns, such multilevel design seems to call for forms of abstraction that bridge between very different quantitative models. For example, what features or patterns in the accelerometer data for a short period of time lead to interesting GPS data around that same time—and vice versa? Such bridging or 'metamodelling' is not yet well understood or established.

Secondly, we explored how the temporal experience of those we observe may play an important role that might not be represented in our current analysis tools. It might be practical to include self-reporting mechanisms that explicitly track user experience of time, to be used to change or adapt representations within temporal data analysis tools. For example, individual players in DoF did not necessarily share the same temporal patterns of engagement in the game, and having a representation of such engagement alongside raw log and video data might helpfully inform how we choose to interpret participants' understandings of the passing of time and the user experience as a whole.

Similarly, such treatment might be used to adapt quantitative algorithmic analysis. One example of such analysis from our own work concerns Bluetooth scan data of the type visualised in Figure 1. By dividing the data set into short windows of time, we can find sets of Bluetooth devices that were within range of each other in a significant proportion of such windows and use this to suggest membership of informal and temporary groups of fans (although further analysis is usually needed to confirm this). Similarly, in other work, we have used an information theoretic measure, mutual information, to find periods of time in which pairs of participants appear to show coordination of movement and system use. For example, in the Treasure game [10], where players competed in teams of two, we used mutual information to examine the extent to which team members coordinated movement as opposed to acting completely independently of their team mate; we could thus automatically find pairs of players who had developed tactics involving separately covering different parts of the game arena. Such techniques generally rely on analysing windows of time of uniform length, but the length of time is chosen rather arbitrarily. Analysts may guess that, for example, windows of 5 minutes' duration would be useful, but these strict divisions of time might not fit well with the way players organise their own reports of the experience. We suggest that algorithms for quantitative data analysis might adapt according to participants' reports of engagement, dynamically adjusting window sizes to fit with, for example, users' impressions of their own activity, the pace of interaction, and the intensity of the user experience.

Conclusion

These two challenges, of wide temporal ranges of data sets and the representation of subjective temporal experiences of participants, thus lead us to question the way that analysis tools might be designed in order to collect together such long-term and fragmented temporal data. As different forms of analysis may require separate tools or components and introduce multiple subjective timelines, we might also consider representing these lengthy subjective time ranges of data in ways that are not restricted by predominantly timeline-based forms. For instance, a single timeline, however augmented, may compress subjective temporal experiences into an objective representation that distorts how a series of events was actually experienced. Whilst not wishing to abandon the timeline representation, we would ask what alternative spatial and representational forms might help us arrange data sets in new ways that are more sensitive to the nature of the user experience. Or, in other words, when is a timeline a useful way of arranging such diverse data, and when does it result in an over-simplifying distortion?

References

[1] Schnädelbach, H. et al. Performing thrill: Designing telemetry systems and spectator interfaces for amusement rides. *Proc. ACM CHI 2008*, pp. 1167-1176.

[2] Sherwood, S. et al. Adapting evaluation to study behaviour in context. *International Journal of Mobile Human Computer Interaction*. In press.

[3] Capra, M. et al. Day of the Figurines Evaluation Report. http://www.pervasive-gaming.org/Deliverables/D12.6-DOF-Evaluation.pdf

[4] Crabtree, A. et al. The cooperative work of gaming: orchestrating a mobile SMS game, 2007, J.CSCW.

[5] Benford, S. and Giannachi, G. Temporal trajectories in shared interactive narratives. *Proc. ACM CHI*. 2008.

[6] Crabtree, A. et al. Supporting Ethnographic Studies of Ubiquitous Computing in the Wild. In Proc. Designing Interactive Systems (DIS), 2006, ACM Press.

[7] Benford, S. et al. Can you see me now? ACM *Transactions on Computer-Human Interaction (ToCHI)*, 2006, ACM Press.

[8] Card, S. K, Mackinlay, J. D., Shneiderman, B., *Readings in information visualization: using vision to think*. Morgan Kaufmann Publishers, 1999.

[9] Morrison, A. et al. Coordinated Visualisation of Video and System Log Data. In *Proc. International Conference on Coordinated & Multiple Views in Exploratory Visualisation*, p. 91-102, 2006.

[10] Barkhuus, L. et al. Picking Pockets on the Lawn: The Development of Tactics and Strategies in a Mobile Game. Proc. *UbiComp*. 2005.

[11] Brundell, P. et al. Digital Replay System (DRS) - a tool for interaction analysis. In Proc. of International Conference on Learning Sciences, Utrecht: ICLS, 2008.

[12] Hancock, P. A. and Weaver, J. L. On time distortion under stress. *Theoretical Issues in Ergonomics Science*:6(2), pp. 193-211, Taylor & Francis, 2005.

Temporal Ecologies in Computing

Michael Terry
HCI Lab, David R. Cheriton School of Computer Science
University of Waterloo, Waterloo, ON, Canada
mterry@cs.uwaterloo.ca

INTRODUCTION

Temporal data highlights change in the midst of stability; reveals trends and patterns over time; and predicts the future as a function of the past. It is of use to designers, researchers, and end users alike, sometimes as an analytical tool, other times as a tool for accomplishing work in and of itself.

For the past year, we have been working with temporal data collected through the ingimp project [12]. ingimp is an instrumented version of the open source GNU Image Manipulation Program (GIMP). The software automatically logs the commands used, basic user interface events (such as use of the keyboard/mouse), and summarizations of the types of documents worked on (e.g., the size of documents). All logged data are sent to the ingimp server where they are available for public analysis. The primary goal of the project is to be able to richly characterize the community's long-term use of the application to help inform future design efforts.

Since its release in 2007, ingimp has been installed over 800 times and has collected over 5,000 log files representing 500,000 commands. This longitudinal data set is the first publicly available data set of its kind, making it a unique resource in the HCI community.

Our work developing and deploying ingimp has revealed a number of issues in analyzing temporal data sets, as well as privacy implications for collection of these data. In particular, we have identified a significant need for visualizations that succinctly summarize application usage over time by thousands of individuals. The project has also highlighted the need to thoroughly consider privacy issues when collecting data over extended periods of time. In particular, data points that are anonymous when sampled once can easily lose their anonymity when sampled repeatedly over time, though this is not always immediately obvious.

The purposes of this paper are threefold. We 1) introduce the notion of *temporal ecologies for computing*, or rich descriptions of human-computer interactions over time, 2) generate a basic design space with which to categorize and consider temporal data in human-computer interactions, and 3) convey challenges inherent in collecting, analyzing, and utilizing these data, based on our experiences with ingimp.

TEMPORAL ECOLOGIES IN COMPUTING

People's relationships with computing technology are ever evolving, even when the technology itself remains static. This evolving relationship with technology can be considered a *temporal ecology* – a description of human-computer interactions over time.

Temporal ecologies enable one to better understand the long-term utility, viability, and broader impact of technology, and can feed directly into design processes, or even inform policy makers. For example, long-term studies of cell phone use in cars have resulted in legislation banning certain uses of cell phones while driving [9].

With the ingimp data set, our goal is to construct temporal ecologies describing individual and community use of the application over time. However, in dealing with the data, we have found a lack of analytical tools tuned to summarizing and visualizing application use by hundreds of individuals. We have also found a need for a vocabulary and framework for structuring our examination and communication of this temporal data. As such, this paper first develops a basic design space for considering temporal data then articulates challenges in working with these data.

TEMPORAL DATA

Temporal data, in its most basic form, is at least a two dimensional quantity: It is a datum with a temporal attribute associated with it. This temporal attribute may be:

- *A timestamp*, or a specific "wall clock" time
- *An elapsed time*, or the difference in time from some prior event or epoch
- *An ordinal value*, or the position of an event in a sequence of events
- *Frequency/rate*, or the number of times an event is observed per unit of time, or the amount of some quantity per unit of time

While not canonical, this list provides some of the most common forms of temporal data in HCI. For example, one can talk about *when* people use a technology during the day (timestamp), the *order* of steps taken when completing a task (ordinal values), the *efficiency* of a user interface design (a rate), or error *rates* (a frequency). One can also talk about changes in any of these data over time (another form of rate). For example, error rates are often measured for users of input devices to measure learning rates of the input device (e.g., [7]).

Related to the concept of ordinal values, one can also talk about *developmental stages*. All living creatures go through well defined developmental stages (e.g., pre-natal, infancy, childhood, adolescence, adulthood); one can similarly consider developmental stages in use of computational technologies. As an example, one could imagine delineating developmental stages related to the adoption of new technology, which might be: initial exploration, focused learning of basic usage, minimal proficiency with the application to perform necessary tasks, reappropriation to new tasks, and so on.

Scales of Temporal Data

Temporal data for computing is of interest in ranges from nanoseconds to years. Real-time, mission-critical systems operate in nanosecond to millisecond ranges, and must interface with users at these scales of time. For example, knowing that the driver wishes to slam on the brakes in a car before actual human motor movement occurs can help the system respond faster, ultimately saving lives. At the other end of the spectrum, observing technological use over months or years enables one to assess broader impacts technology, as cell phone use in cars illustrates.

Sources of Temporal Data

Given the basic types of temporal data, one can also consider sources of temporal data. In practice, there are an infinite number of such sources – anything that can be measured can be measured over time. Accordingly, we divide this space to consider three independent sources relevant to the HCI community, namely, the system, the user, and the interactions between the two.

From the system's perspective, one can consider data such as:

- The system's hardware and software configurations, including what hardware is available, what software is installed, how items are configured, network connectivity, BIOS versions, etc.
- System "health," such as whether the system's hardware is operating correctly, whether it is infected with viruses, whether it is being attacked, etc.

From the user's perspective, one can consider:

- Goals, desires, and preferences
- Cognitive and physical abilities, and biometric data
- Environment and context

Finally, one can consider interactions between the user and system:

- What the user does with the system, at both macro and micro scales, such as what tasks it is used for, the individual actions performed, etc.
- Users' knowledge of, and proficiency with, the system, including error and performance measures
- Overall satisfaction, aesthetic appeal, usability, utility, and other qualitative assessments of the system and interactions with the system

Combining these data with temporal dimensions yields measures such as:

- Learning, error, and productivity/performance rates
- How the user adapts to, adopts, and reappropriates the system
- Software/feature/hardware uptake, or how quickly users adopt new hardware, software, and its features
- Mental and physical fatigue related to use of the system
- Aesthetic and design hysteresis, or whether the aesthetics and design of the system retain their appeal and utility over time, as the user's use and perceptions of the system evolves
- Data useful for security forensics, such as how a system became compromised (e.g., through a successful phishing attack)

Again, this list represents only a sampling of the types of temporal data possible in examining human-computer interactions, but provides a starting point for discussions on the subject.

Utility of the Data

Temporal data are useful to designers and researchers alike, where they can serve diagnostic, analytical, and predictive purposes. They also have direct utility as tools for *end-users*. For example, the recording of *past states* enables:

- Error recovery via undo, source control, and backups
- Experimentation, for example, by using undo to return to a previous state after experimenting with the application or alternative solution possibilities
- Refinement of past actions through mechanisms such as selective undo [1] or editable graphical histories [6], both of which allow one to tweak past actions
- Illustrating a history of edits to understand how a design was generated. Examples include visualizations of past actions as storyboards [11] or comic-like panels [6]

- Augmenting the user's memory. This augmentation can be used to enhance cognition during routine work tasks (e.g., Rhodes' Rememberance Agent [10]) or to assist those with cognitive disabilities (e.g., Microsoft's My Life Bits project [4, 8])
- Task automation. Remembering what was entered in data fields in the past enables that same data to be used in the future. Making available the most recently used commands on the command line has also been shown to be highly effective at increasing efficiency [2]

Learning Systems

Temporal data can also be used to learn systems, such as adaptive and predictive systems. Examples include menus that adapt to show only the most frequently used items, Hurst's magnet-like menus that make selecting previously used items easier [5], and Amazon.com's product suggestions based on past purchases.

These dimensions – types of temporal data, sources of temporal data, and uses of the data – serve to sketch a basic design space for temporal data. From this basic design space, we now turn to a set of challenges we have identified in dealing with temporal data in the ingimp project.

TEMPORAL DATA: CHALLENGES

As mentioned, our interest in temporal data stems from our desire to characterize the long-term practices of the GIMP community. In this section, we focus on two challenges encountered in working with these data: Analysis and privacy issues.

Challenges in Analyzing Temporal Data

While instrumentation of applications is a common practice, few standards or tools exist to assist in the collection, analysis, and summarization of temporal data in interactive environments. Arguably, the largest selection of tools available in this area resides in the domain of web applications and website usage. However, web applications have comparatively less functionality than rich client content-creation applications, the latter of which boast hundreds of commands available to the user at any point in time. Summarizing the activity of thousands of users over thousands of sessions remains a significant challenge for such applications, as we describe.

Visualizing Temporal Data

With the ingimp data set, we would like to be able to generate at-a-glance visualizations that characterize the types of tasks users perform, their proficiency in those tasks, the types of documents they operate on, how their documents change over time, and how they utilize the application in larger workflows. However, the qualitative data collected thus far hints at the challenges of creating such at-a-glance visualizations. Based on users' entered "Activity Tags" (free-form descriptions of how they plan on using the software at start-up), users indicate they use the software for gel analysis, map making, icon editing, avatar texturing, cell shading, and videogame graphics, just to name a few. Across these tasks, there is considerable overlap in the functionality used (for example, cropping an image), but at the time, great variety in how and why the functionality is used. Ideally, we would like to be able to look out across this sea of users and get a feel for this variety and similarity, much in the same way one can take a snapshot of a cityscape and gain a sense of the city, its landscape, and its diversity.

Current tools for visualizing software use include traditional graphs (line graphs, bar charts, histograms) and the more recent Chromogram [13], the latter of which uses hue to convey what commands are used, over time. For our work with ingimp, we have also created a basic summarization of individual usage whereby a stick figure persona is generated, embellished to reflect the user's usage of the software (Figure 1). All of these techniques provide basic summary information, but more work is necessary to condense rich, multidimensional, temporal data into compact visualizations tuned to conveying software usage.

We are also interested in *dynamic* presentations of the community's use of the application, including time lapse equivalents of viewing application usage by the community. Time lapse is an attractive representational form since it collapses long expanses of time into more human-consumable scales. However, the challenge in this domain is finding representational forms that adequately display activity flows over time. Replaying screenshots or video at sped-up rates is one possibility, but we have not found this to be as elegant or as useful as traditional time lapse video of natural events. Instead, it is likely that it is necessary to synthesize new renderings of interaction to make this technique viable.

In addition to compressing long spans of time, it can be equally important to unpack very short segments of time. As an example, our data indicate that few users read the software's text-based consent agreement prior to using ingimp. Accordingly, we are redesigning the agreement to increase the chance that users actually take notice and read it. In this case, we need to examine rapid eye movements to understand how the redesigns capture and retain attention. Examining interactions at these extremely small scales of time is rare in HCI, with the majority of work being conducted in areas such as motion kinematics or human factors.

Querying Temporal Data

The ability to interactively query temporal data is essential to constructing temporal ecologies of use. Currently, the state of the art in querying temporal data amounts to specifying ranges of time, or searching for particular events. However, one can also imagine specifying patterns, trends, and anomalies as query "terms" when dynamically querying the data. Some experimental environments allow one to specify patterns using awk-like languages [3], but such capabilities are rarely available. While mature techniques *do* exist within the artificial intelligence and machine learning communities for extracting these types of features from temporal data (e.g., Hidden Markov Models (HMMs), particle filtering, and artificial neural networks), few tools have crossed over into the HCI community for interactively exploring temporal data along these dimensions.

Privacy Issues in Using Temporal Data

Our work in creating and deploying ingimp has highlighted the privacy issues that can arise when collecting data over time. In particular, a single data point that is sufficiently anonymous on its own can easily lose its veil of anonymity when collected over time. As an example, ingimp records each event with millisecond accuracy. Each press of a key is recorded, though the actual key pressed is not recorded. However, over time it is conceivable that one may be able to infer likely keystrokes by examining the time *between* each keystroke. Accordingly, we are currently modifying the logging of keystrokes to aggregate the events, rather than reporting each single keystroke. As this example highlights, the introduction of temporal qualities to data can render single-point anonymous data non-anonymous in ways that are not necessarily obvious at first blush. Thus, guidelines for privacy-preserving ways to collect temporal data could be of use to this community.

SUMMARY

This paper has sketched out a basic design space for considering temporal data in human-computer interactions and has identified some noteworthy challenges in utilizing and collecting these data. Within this workshop, our primary interests lie in contributing to best practices in visualizing, summarizing, and analyzing temporal data of human-computer interactions.

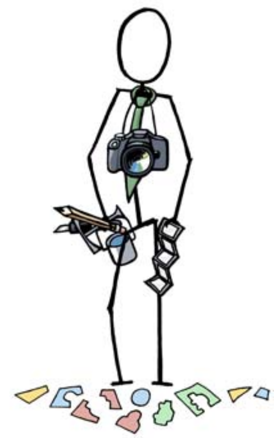

Figure 1. The ingimp website dynamically generates *personas* for each user. An ingimp persona summarizes the most frequently used commands (held in the right hand), the typical size of images worked on (held in the left hand), and the types of tasks (indicated here by the camera and clippings of images on the ground).

REFERENCES

1. Berlage, T. A selective undo mechanism for graphical user interfaces based on command objects. *ACM Trans. Comput.-Hum. Interact.* 1, 3 (Sep. 1994), 269-294.
2. Greenberg, S. and Witten, I. H. How users repeat their actions on computers: Principles for design of history. In *Proc. ACM SIGCHI Conference on Human Factors in Computing Systems* (ACM CHI'88). pp. 171-178
3. Hilbert, D. M. and Redmiles, D. F. 2000. Extracting usability information from user interface events. *ACM Comput. Surv.* 32, 4 (Dec. 2000), 384-421.
4. Hodges, *et al*. Wood: SenseCam: A Retrospective Memory Aid. In *Ubicomp 2006*: 177-193.
5. Hurst, A., Mankoff, J., Dey, A., and Hudson, S.E. Dirty Desktops: Using a Patina of Magnetic Mouse Dust To Make Common Interactor Targets Easier To Select. In Proceedings of *UIST '07*.
6. Kurlander, D. and Feiner, S. Editable graphical histories. In *Proceedings of the IEEE Workshop on Visual Languages* '98. pp. 127-134
7. Lyons, K., Gane, B., and Starner, T. Experimental Evaluations of the Twiddler One-Handed Chording Mobile Keyboard. *Human-Computer Interaction*. Volume 21(4), pp. 343--392, 2006
8. Mihailidis *et al*. The use of an intelligent prompting system for people with dementia. *Interactions* 14(4): 34-37 (2007)
9. Redelmeier DA, Tibshirani RJ. Association between cellular-telephone calls and motor vehicle collisions. *N Engl J Med* 1997;336:453-8
10. Rhodes, B., and Starner, T. The Remembrance Agent: A continuously running automated information retrieval system. In *Proceedings of The Practical Application of Intelligent Agents and Multi Agent Technology* (PAAM '96), London, UK, April 1996, pp. 487-495.
11. Su, Sara. Visualizing, Editing, and Inferring Structure in 2D Graphics. Doctoral Symposium, ACM Symposium on User Interface Software and Technology (UIST), October 2007
12. Terry, M., *et al*. ingimp: introducing instrumentation to an end-user open source application. In *Proceedings of CHI '08*. pp. 607-616
13. Wattenberg, M., Viégas, F, Hollenbach, K. Visualizing Activity on Wikipedia with Chromograms. In *Proceedings of Interact, 2007*

"Best If Used By": Expiration Dates for Email

Manas Tungare
Dept. of Computer Science, Virginia Tech.
Blacksburg, VA, USA.
manas@tungare.name

Manuel A. Pérez-Quiñones
Dept. of Computer Science, Virginia Tech.
Blacksburg, VA, USA.
perez@cs.vt.edu

Abstract
We recognize the ephemerality of certain kinds of email received, and propose the use of an expiration date tag to indicate its lifetime. We hypothesize that the use of such a tag will assist personal information management (PIM) by providing users the ability to prune their email archives automatically, and take other actions as appropriate. We situate our proposal of expiration tags within the current PIM literature, focusing on the research problems they may help solve. We conclude with a discussion of how expiration tags can be set, retrieved, and acted upon by mail clients.

Keywords
Personal information management, email, expiration dates, temporal information.

ACM Classification Keywords
H.5.2 Information Interfaces and Presentation: User interfaces – Evaluation/ methodology

CHI 2009, April 4 – 9, 2009, Boston, MA, USA
ACM 978-1-60558-246-7/09/04.

Introduction
The past few years have seen an unprecedented rise in the use of email for communication, collaboration, information management and several other tasks it was not explicitly designed to perform [9, 3]. The cognitive costs associated with the manual filing and pruning of information archives overwhelm many users [9]. Messages vary in relevance, importance, timeliness and attentional requirements. Several strands of research have examined how such properties of email can be computed and used to provide a better user experience [3]. Social network analysis has been used to help triage incoming email as well [4]. The time-sensitive aspects of email have been recognized [6], but research has focused on prospective task management rather than retrospective archive management.

Examples
We conducted a quick and informal analysis of our inboxes and those of colleagues in our research group. We found an incredible variety in how long the information content within email messages found in the inboxes stays valid, useful, or pertinent. Here are some examples:

- **IM-style 1:1 communication.** *(extremely short term)* One-line (or even one-word) messages are often sent in lieu of a phone call: e.g. "Lunch? ", "Running late to meeting", "Movie tonight? ".
- **Awareness notifications.** *(extremely short term)* Several web-based services send notification alerts to users when certain monitoring criteria are met. User action may or may not be solicited or expected. E.g. "bill is due in 2 days", "X added you

as a friend", "your order was received", "your package has shipped", "free donuts in break room".

- **Project-related communication.** *(short term)* Email related to an ongoing project will soon be outdated after the project is completed. E.g. "Draft 5 attached", "I booked my tickets for Monday".
- **Discussions with archival value.** *(medium term)* E.g. research ideas or conversations related to specific technical problems that may be required later. Although it may be unclear when this information may be needed, it is clearly important to archive it.
- **Affective conversations.** *(long term)* Conversations with significant others or immediate family may be saved for their nostalgic future value.

The different types of emails shown in the previous listing have different implications for how users manage their email. Email clients fail to take into account the distinction among these types of messages, thus lending inadequate support to information management.

Relevant Prior Work

Files have temporal properties; Barreau and Nardi [1] classified files as ephemeral, working, or archived data. Gwizdka [5] proposed the classification of email into four types: prospective, ephemeral, working and retrospective. However, neither of these approaches have led to the development of solutions towards harnessing the time dimension to assist users in information management.

The Keeping Problem and Post-Valued Recall in PIM

Email is no longer just a communication medium; it also serves an archival role. In PIM, the "keeping decision" [8] refers to the choice a user must make about whether a particular information item is worth keeping for potential later lookup. However, when an email is received, it is not immediately clear how long the message will continue to be useful. Post-Valued Recall (PVR) [10] refers to the interest a user may have in recalling information whose value is not recognized until some time after its initial retrieval. Since users cannot decide right away what to do with an email, they let it linger [11]. However, seldom do users go back to these messages to clean them up later.

Immediate Filing is not very practical

Email is managed in different ways by different people; Gwizdka suggests [7] that handling incoming information immediately is an ideal case, but it is not practical for several reasons. The cost of a search multiplied by the probability that a particular information may be searched for, is much less than the cost of constantly having to file, tag, and sort email. Filing is a cognitively difficult task; while some users are natural cleaners, others are keepers. The upshot is that email continues to stay in inboxes longer than necessary. There is another problem with immediate filing, as identified by Whittaker and Sidner [11]: once an email is filed away, it is less available to remind the user about that topic (less chance of opportunistic reminding).

Life archives

As the information haystack grows larger [2], it becomes harder to find the proverbial needle. It is important that non-essential irrelevant information be pruned from an archive as early as possible (though no such decision can ever be taken with 100% accuracy [8].)

Setting an Expiration Tag

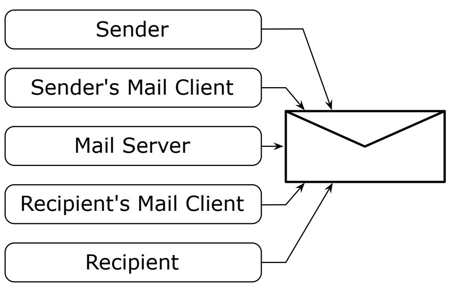

Actions taken on Expiration Tags

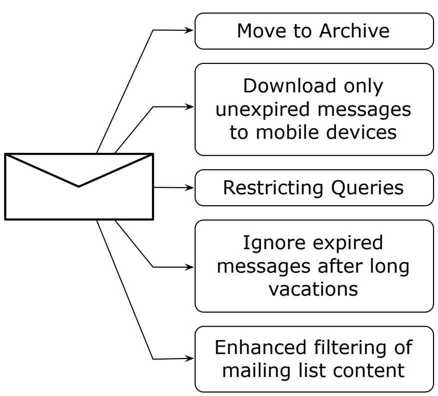

A Solution: Expiration Date Tags

An Expiration Date Tag is an email header that provides a best estimate of when an email message is projected to be irrelevant to the recipient. It serves as an indication of its time sensitivity — an attribute that is not captured by any existing headers. A complete description of the syntactic aspects of expiration date tags is beyond the scope of this paper. We refer you to a deeper discussion online [1].

Applying an expiration date splits the task of filing emails into two independent subtasks: an expression of intentionality and the performance of the action. A user's intention to file an email can be expressed as soon as a message is received (which is ideal according to Gwizdka [7]). The filing and archiving itself, however, is done at a later time automatically (which avoids the problem of lack of opportunistic reminding noted by Whittaker [11]). In addition, support for expiration dates provides an opportunity for various entities (other than the primary recipient) to apply the tags automatically. Specifically, it supports the management of prospective information in email [6].

Setting Expiration Date Tags

Expiration Dates can be set by several entities, not just the primary user. They could be set:

- By the sender of an email who can make a reasonable assumption of the relevance of her email to the intended recipient; e.g. credit card payment reminders can be automatically set to expire 5 days after the due date.
- By the mail server software that intelligently tags email based on common patterns seen across multiple users (like spam filters do);
- By the recipient's email client, based on heuristics; (say, if a pattern has been observed that emails with certain subject lines are deleted by the user in X days)
- By the recipient's email client, based on a user-defined rule set; ("from:notifications@facebook.com ⇒ expire in 5 days")
- Or explicitly by the recipient in a spring cleaning session.

The simplicity and flexibility of the tag means that any party involved in the transmission of the email can modify/update it. Security concerns about potential tampering by men-in-the-middle may be assuaged by knowing that this is no more vulnerable than the rest of the email.

Acting Upon Expiration Date Tags

An expired email need not (and should not) immediately be deleted if the user does not so desire. It is indicative, not prescriptive. Here are some ways we expect an expiration tag to assist in personal information management:

- An email past its expiration date could be automatically moved from the Inbox to Archived items;
- Expiration tags can be used in complicated searches by restricting a query based on expiration dates (e.g. show all emails that are due to expire in

[1] http://manas.tungare.name/blog/email-should-have-expiration-dates/

the next week);

- Due to resource limitations of mobile devices, typically only the most recent few emails are downloaded and displayed. Incorporating expiration dates into the decision can cause more relevant messages to be shown, while expired messages stay hidden.
- Automatic pruning can be performed during spring cleaning sessions; this provides the user a quick option to delete expired messages permanently without having to deal with each individually.
- Expiration can greatly reduce the deluge of pending email after a vacation. Notices about free donuts and missed meetings can automatically be removed by the system.
- Attaching expiration dates to messages sent to distribution lists provides a new alternative to the four strategies outlined by Mackay in [9].

Future Work

We are building a prototype of a system that can read and write expiration tags. Tags will be applied in two ways: by an IMAP client process that polls and processes a user's inbox, and instrumented email clients (Apple Mail.app, Microsoft Outlook). We look forward to sharing preliminary results within a few months.

References

[1] D. Barreau and B. A. Nardi. Finding and reminding: file organization from the desktop. SIGCHI Bull., 27(3):39–43, 1995.

[2] M. Czerwinski, D. W. Gage, J. Gemmell, C. C. Marshall, M. A. Pérez-Quiñones, M. M. Skeels, and T. Catarci. Digital memories in an era of ubiquitous computing and abundant storage. Communications of the Association for Computing Machinery (CACM), 49(1):44–50, 2006.

[3] N. Ducheneaut and V. Bellotti. E-mail as habitat: an exploration of embedded personal information management. interactions, 8(5):30–38, 2001.

[4] D. Fisher, B. Hogan, A. Brush, M. Smith, and A. Jacobs. Using social sorting to enhance email management. In HCI Consortium Meeting, February 2006. HCI Consortium, 2006.

[5] J. Gwizdka. Timely reminders: a case study of temporal guidance in pim and email tools usage. In CHI '00: CHI '00 extended abstracts on Human factors in computing systems, pages 163–164, New York, NY, USA, 2000. ACM Press.

[6] J. Gwizdka. Supporting prospective information in email. In CHI '01: CHI '01 extended abstracts on Human factors in computing systems, pages 135–136, New York, NY, USA, 2001. ACM.

[7] J. Gwizdka. Email task management styles: the cleaners and the keepers. In CHI '04: CHI '04 extended abstracts on Human factors in computing systems, pages 1235–1238, New York, NY, USA, 2004. ACM Press.

[8] W. P. Jones. Finders, keepers? the present and future perfect in support of personal information management. First Monday, 9(3), 2004.

[9] W. E. Mackay. Diversity in the use of electronic mail: a preliminary inquiry. ACM Trans. Inf. Syst., 6(4):380–397, 1988.

[10] J. Wen. Post-valued recall web pages: User disorientation hits the big time. IT & Society, 1(3):184–194, 2003.

[11] S. Whittaker and C. Sidner. Email overload: exploring personal information management of email. In CHI '96: Proceedings of the SIGCHI conference on Human factors in computing systems, pages 276–283, New York, NY, USA, 1996. ACM Press.

Just for me: Personal Applications of Life Tracking and Activity Capture

Max Van Kleek
MIT CSAIL
32 Vassar St
Cambridge, MA 02139 USA
emax@csail.mit.edu

David R. Karger
MIT CSAIL
32 Vassar St.
Cambridge, MA 02139 USA
karger@mit.edu

Brennan Moore
ZAMIANG
126 Binney St.
Cambridge, MA 02139 USA
brennan@zamiang.com

mc schraefel
School of Electronics and
Computer Science
University of Southampton
mc+chi@ecs.soton.ac.uk

CHI 2009, April 4 – 9, 2009, Boston, MA, USA
ACM 978-1-60558-246-7/09/04.

Abstract
In the past few years, the explosion of interest in social sharing on the Web has made popular the practice of publishing one's everyday life activities on the Web. The resulting availability of rich, detailed data about an individual's activities presents newfound opportunities for personal applications beyond social consumption. In this paper, we describe several systems which make use of temporal personal activity data aggregated from Web sites for several personal purposes: to help users more easily re-find their notes, automate routine reactive tasks, organize their digital resources and reflect on their personal life metrics. Each of these applications use aggregated temporal data in different ways, providing different perspectives of how temporal activity data can support interaction.

Introduction
This excitement of the social sharing on "Web 2.0" has driven the creation of a huge array of web sites and tools make of the chronicling of the minutiae of everyday life activities easy and fun. The result is that social sharing on the Web has spawned an unprecedented quantity of detailed, real-time information about people's lives. While the social implications and potential for such sharing are still being explored, we believe that this data is also useful towards solving interaction problems closer to home –

personal applications for the individual. In particular, we examine the opportunity for high-fidelity records of user activity gleaned from the Web to support the individual in the following three ways:

- *as a memory prosthesis*, e.g., to support the recall and memory of past events and information;
- *to support simple task delegation*, e.g., the construction of reactive tasks that respond to the individual's activity and save the user from having to perform these actions manually;
- and to increase the individual's self-awareness, by making salient long-term statistics of the individual's life through longitudinal analysis of the user's activities. These metrics could include time and attention spent on projects, social contact with individuals, and tracking physical life activities such as amount of exercise, sleep and nutrition.

We briefly describe applications that use temporal actiity data to fulfill each of these forms of support in the remaining sections of the paper.

Memory support for personal information-keeping

Individuals keep track of an immense quantity of information that they rely upon to effectively perform everyday tasks and fulfill daily responsibilities in work, social and personal contexts each day. This information might pertaining to reminders of tasks to be done, important phone numbers, web sites, books to read, birthdays, wish-lists, instructions on how to do things, et cetera. We examined how users create and maintain their information at work and home using current tools (both paper and software) in an artifact and interview study [1], which revealed that there are many kinds of information that people frequently need that "don't fit" or are otherwise burdensome to manage using conventional personal information management tools. For these kinds of information, people often resort to using e-mail (i.e., e-mailing themselves), or stashing information in "misc" files, folders, and post-it notes on their desks/workspaces. While these strategies result in the storage of information that affords some of the informational needs surrounding these information items (such as visibility, in the case of a post-it note, or availability in the case of e-mail), information kept using these strategies must be manually managed, and lack the facilities that more structured PIM tools readily provide: reminder alarms, multiple views (calendar/todo), and automatic filing/organization.

Therefore, we sought to develop a tool that would afford the flexible and low-effort storage and creation of information afforded by note-taking tools, but that would also support reminding and use once it was captured. Our systems, Jourknow and its successor list.it [4], apply user activity data (e.g., what the user is doing, where, and when) aggregated from the Web to learn how relevant particular notes are to various contexts/activities of people's lives. Further relevance information is computed from features extracted from the note, specifically dates, times and names of people and places and web sites, enabling the specification of "context-sensitive alarms". Finally, these relevances are then used to automatically display notes that are likely to be useful to the individual as she goes about her daily activities. Both of these systems display these learned associations to the individual, allowing them to inspect and modify these associations and improve reminding behavior. A screen shot is visible in Figure 1.

Activity-reactive personal automation

The second application surrounds ways to relieve the user from the attentional burden of simple activity-reactive tasks. Our approach, embodied in a system called **atomsmasher** lets users specify their desired autonomous or reactive autonomous behaviors directly, by "programming" a simple reactive script to perform tasks for them. In the same way that calendar alarm/reminders and e-mail filtering functions relieve the user from paying attention to the time and each incoming email respectively, these scripts can perform an action or combination of actions in response to an event about user activity arriving from various Web streamss. For example, AM can be used to easily set up an adaptive "Away responder" that can automatically determine, based on a user's calendar entries or location, when and how to automatically reply to messages of particular types. Atomsmasher users specify behaviors using a simplified constrained natural language input interface to specify behaviors in terms of "when X happens do Y". More information about atomsmasher can be found in [1].

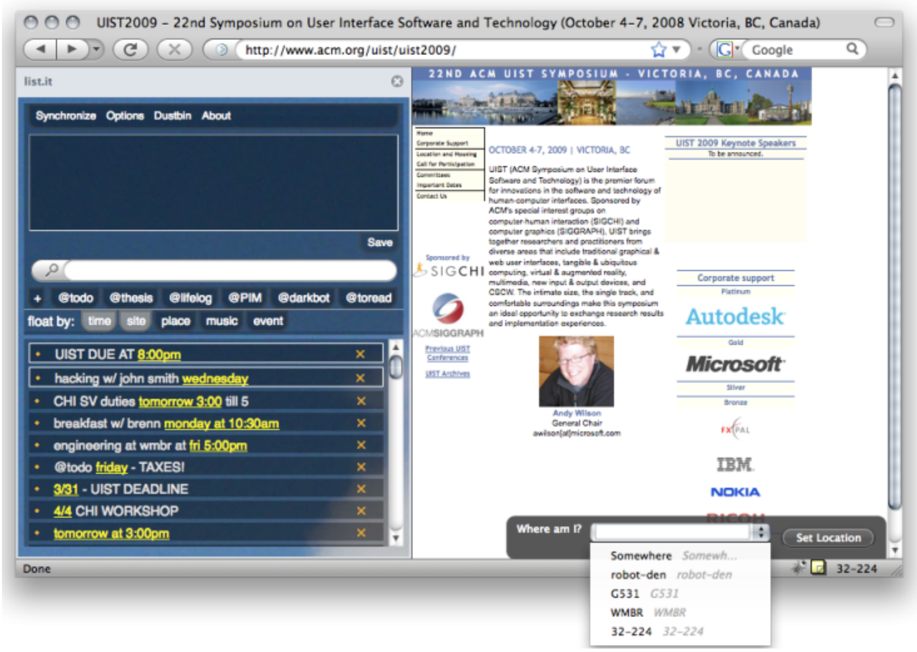

Figure 1 - List.It with Notes that Float (NTF), which supports multiple "floation modes" by content features (date/time/entities) and activity correlations.

Self-awareness and life management

Our third application seeks to let people easily reflect upon their personal information in context of their life activities. Our approach unifies three formerly separate ideas -- temporal and activity based organization of personal documents (proposed by Lifestreams [2] and [1]) and statistical visualization of life activities (introduced by services such as DAYTUM[1] and mycroscosm[2]). While the former two approaches use time and user-defined activities as organizational principles for document collections, the latter two "micro-tracking" sites facilitate automatic collection and publishing of select summaries of their life activities. Our perspective is that these systems could greatly benefit from being integrated -- enabling summaries of life activities to contextualize personal information spaces, and vice versa-- allowing personal information spaces to encourage reflection on self activity.

---

[1] http://www.daytum.com

[2] http://mycrocosm.media.mit.edu

To this end, we are building a life tracking service based on our user modeling framework, that provides integrated views of heterogeneously sensed activities and contexts, naturally situating interactions the user has with informational resources and people at the proper locations within these views. Although, like in Lifestreams, time the main organizing axis, instead of supporting only linear timelines, we wish to explore the use of periodic timelines that can "fold" the user's past into a single view to facilitate reflection further into the past by days, weeks or months. These views aim to make salient regular patterns (such as temporal periodicities) in certain activities, in addition to showing trend of particular activity summaries, such as increased or decreased contact with particular individuals, physical activity, nutrition, sleep, etc. We imagine that such trends might help individuals to explain/better understand their states of well-being. An early design mockup is visible in Figure 2.

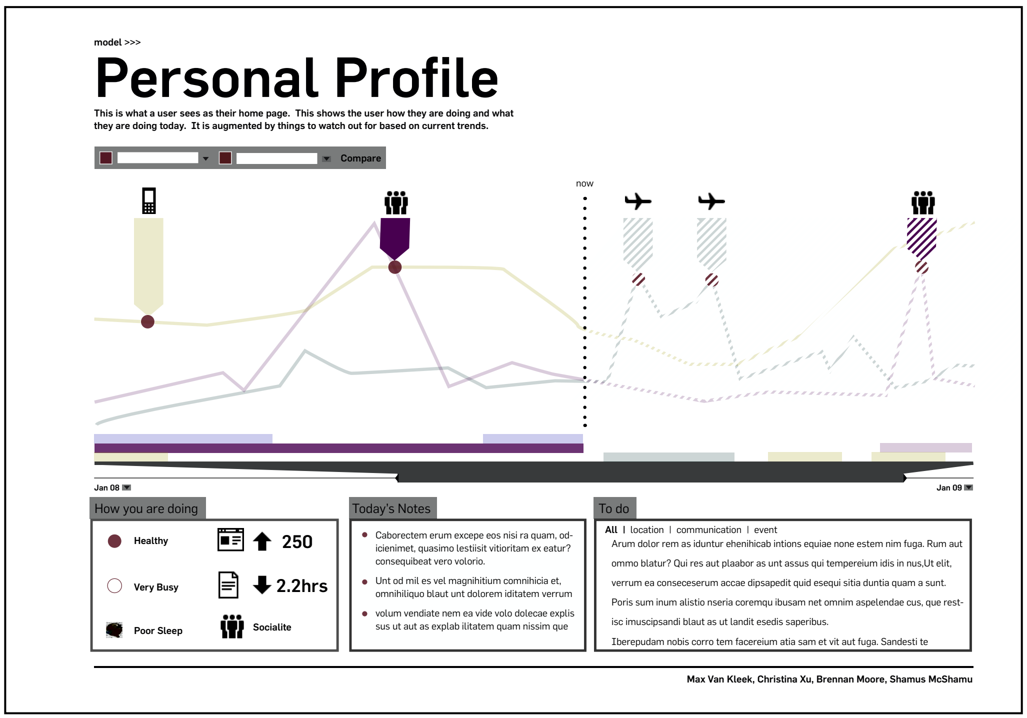

Figure 2 - Eyememine, the "activity desktop", which contextualizes personal information items in captured life activity data, to encourage self-reflection and facilitate organization.

Acknowledgements

The following individuals have contributed to the work described: Greg Vargas and Katrina Panovich (list.it), Jamey Hicks (list.it and Organic Indoor Location), Michael Bernstein (early PLUM and Jourknow), Paul André (Atomsmasher), and Brennan Moore (Eyememine). This work was supported by Nokia Research, The Web Science Research Institute, and the National Science Foundation.

References

[1] Bardram, J. "Activity-based computing: support for mobility and collaboration in ubiquitous computing" *Personal and Ubiquitous Computing* v9, issue 5, pp.312-222, Springer, 2005.

[2] Freeman, Eric T. "The Lifestreams Software Architecture," Ph.D. Dissertation, Yale University Department of Computer Science, May 1997.

[3] Van Kleek, M., P. André, D. Karger, m.c. schraefel. "Mixing the reactive with the personal: Opportunities for end user programming in Personal information management", to appear in *End User Programming for the Web*, A CHI Workshop, 2009.

[4] Van Kleek, M., K. Panovich, G. Vargas, D. Karger, m.c. schraefel, "Notes that Float: Applying content and context features to facilitate flexible reminding in notes-to-self". In submission to UIST 2009.

Mirror bodily experiences over time

Elsa Kosmack Vaara
Swedish Institute of Computer Science
Box 1263
SE-16429 Kista, Sweden
+46 8 633 1500
elsa@sics.se

Kristina höök
Swedish Institute of Computer Science
Box 1263
SE-16429 Kista, Sweden
kis@sics.se

Jakob Tholander
Swedish Institute of Computer Science
Box 1263
SE-16429 Kista, Sweden
Jakob.tholander@sh.se

CHI 2009, April 4 – 9, 2009, Boston, MA, USA
ACM 978-1-60558-246-7/09/04.

Abstract
The Affective Health system is a mobile lifestyle application that aims to empower people to reflect on their lives and lifestyles. The system logs a mixture of biosensor-data and other contextually oriented data and transforms these to a colorful, animated expression on their mobiles. It becomes a mirror of aspects of their everyday activities, empowering users to see patterns relating to stress. People in different cultures or people with different physiological and psychological attributes have quite different perceptions and associations of time. We explore the time dimension of our system through working through a set of different designs that organize events as time going linearly forward, in a circular movement or relating to geographical places. Here we discuss our design process and the problem of presenting data from only one temporal perspective.

Keywords
Visualization, graphics, interaction ,time

ACM Classification Keywords
H5.m. Information interfaces and presentation

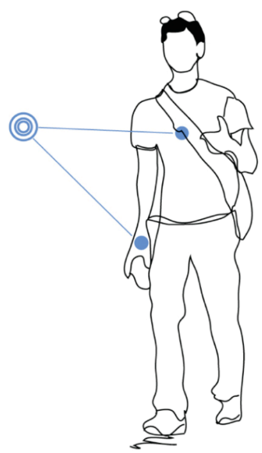

The bodily data is collected via biosensors worn by the user.

The Affective Health application
The Affective Health system is a mobile application that aims to make people equipped to better cope with stress and stressful situations. It is a tool for visualizing patterns and trends of bodily and contextual information over time. This information is to empower users to reflect on their behavior in everyday life.

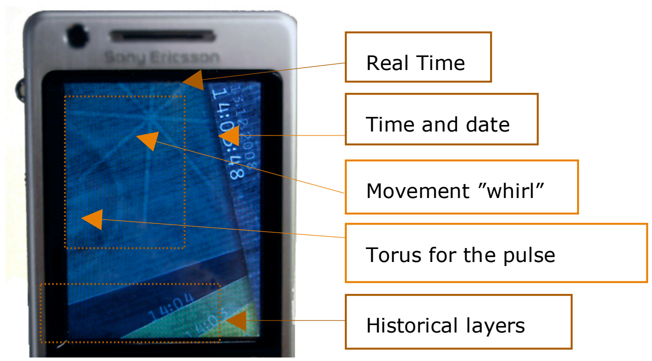

Fig 1. The current application running on a mobile phone.

At the beginning of the project we identified two kinds of events that can give the users valuable information to reflect on: data about bodily reactions and contextual information that connects the bodily data to the ongoing everyday life. The system does not provide a diagnosis of stress. It is well known that people's ability to deal with stress is connected to their subjective experience of their ability to deal with stressful situations (REFS).

We don’t believe that the data will give users an entirely true or full story about their everyday behaviors, stress reactions or life styles, but the material can let users link the data to their subjective experiences. It can support people in forming their own stories about their own lives and in their own ways.

The bodily data is collected via biosensors worn by the user. The sensor signals are transmitted via a Bluetooth device to the application running on a mobile phone. Depending on reliability, wearability and maintenance requirements [1] we chose to use three kinds of sensor data: we monitor some aspects of users' arousal through GSR (Galvanic Skin Response), users' pulse through a heart rate sensor since prolonged increases in pulse may also indicate stressful experiences. Finally, to help users to distinguish stress from normal physical activity, we also measure movement through accelerometers.

Some stress medicine experts claim that physical exercise that evoke similar processes in the body as negative stress does, can train our bodies to be able to deal with stress in a purely biological way [2]. If we go out jogging, our sweat levels and pulse will increase. The jogging experience is probably a positive experience, apart from the negative stress reaction that evokes similar bodily reactions. Apart from movement, we also collect other contextual data that can help users to remember and distinguish different activities in their lives assisting them to sort out what is stressful in a negative way. We collect various data from their mobile phones, such as photos, text messages sent and received, Bluetooth presence of other devices nearby, etc.

Related products
There are a number of body monitoring systems (e.g emWave and the CocoroMeter) [3] developed that are claimed to have the ability to monitor emotions and stress through bodily signals such as heart rate, temperature, movement, hormonal levels of e.g. cortisol or adrenalin,

arousal. Many of these systems though, do not help users connect the sensor information to everyday behavior. The overall approaches are more about measuring users' stress levels and provide this feedback as diagnosis and warning. [4]

Research challenges

Peoples' activities and experiences occur in time and time is a thread to which information can be tied. People in different cultures or people with different physiological and psychological attributes have quite different perceptions and associations of time. This can become a problem and a source for stress related issues since the organization of our post-industrial society is tied to a certain perception of time.

How can we through Affective Health present alternatives to the linear procession of past, present, future? How can we inspire people to create meaning from the bodily and contextual information that we present? The ultimate goal with Affective Health is to make the user prioritize reflection to reach insight in her/his behavior and the reasons for it. Even though we assume that users of this application are motivated to get to know themselves better, it is important to have a desirable and beautiful application to inspire usage.

It is not possible to represent our continually changing lives with one complete and accurate picture. In the Affective Health project we assemble a collage of information that creates living pictures and impressions. This builds a dynamic model to reflect upon and compare to different events and occasions in life. To better understand how to represent stressful experiences back to users, organized in some kind of temporal order, we decided to explore the problem through creating different design concepts where temporality took on different appearances. [fig 1]

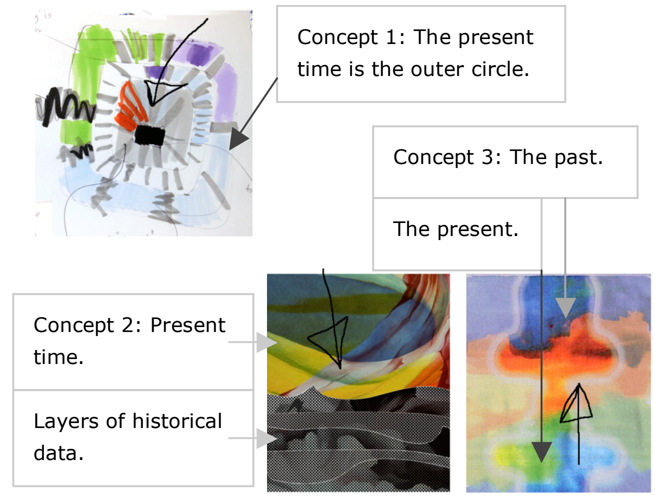

Fig 1. Sketches of the initial three design concepts showing temporal progressing in different forms.

The first concept was built on the idea of pulse and movement vanishing into a time-tunnel where one could trace the past at the end of the tunnel. Our second design concept was inspired by nature and the idea of leaves falling to the ground, building layers of historical data. A third concept showed time as a colored wave where the bottom of the wave represents what is happening right now and the past is floating upwards. In all these design concepts we represented arousal using color. This is based on Anna Ståhl's design work [5] where she uses colors to represent emotional arousal - closely related to what we aim to show here.

Fig 2. Test with Lo-fi paper model.

During the ongoing design process we continually collect feedback from users in different ways, e.g. Lo-fi prototyping. When choosing concept to work in depth with we involved users by presenting 2D paper-models of the basic functions: time, history, heart rate, physical activity, and arousal [Fig 2]. Six unprepared participants were asked to share their understanding and associations of the models.

In this way we aimed at getting a direct reaction to how the different concepts could reflect their experiences. We chose to continue working in depth with the "nature" concept since this, in the early user studies proved to evoke most associations with time. We got reflective comments from people who associated the model with e.g. a "river of time" and "the ground". This corresponded quite well with the thoughts we had about the concept.

One of the main ideas of Affective Health is for the user to be able to browse back in history to explore patterns and characteristics in their lives and behaviors. In the current implementation of the "nature" concept running on a mobile phone, users can scroll back in time divided into standard units such as in minutes, hours, weeks etc. This is a starting point of a temporal concept, but this solution has turned out to be far to narrow.

The problem of visualizing time

Time is subjective and depending on how we live our lives, it is perceived differently. In the current version of Affective Health we see only one perspective of time; the linear. The linear perspective, by many people understood as something that goes by, (the past-present-future idea) is the one characterizing our culture and perception of time today but our perception of time changes continually depending on our activities and memories: You can for example live a slow Saturday afternoon with tea and scones, or fight towards time a busy Monday afternoon.

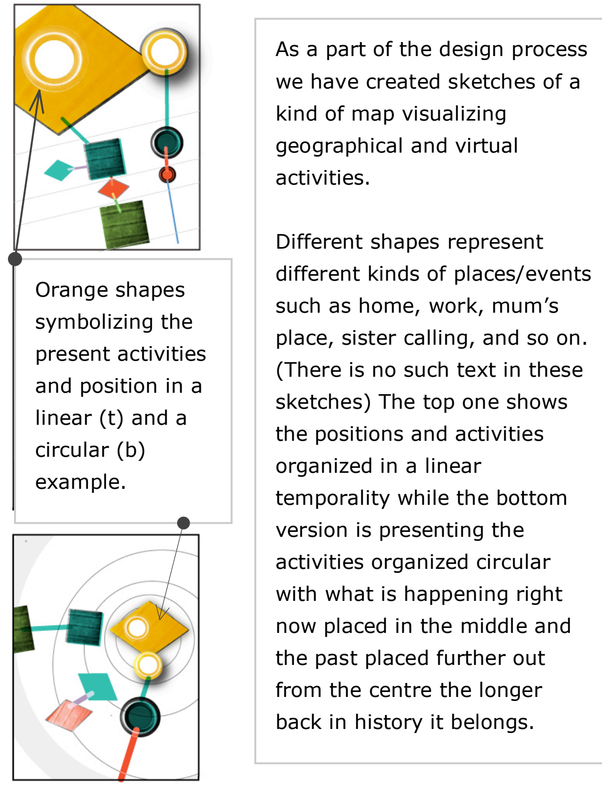

Fig 3. Two different presentations of temporal activities and geographical positions.

This has to be mirrored in the interface of Affective Health and the way we believe we could change this is to not follow just one perception of time, but add several perspectives to give a possibility for a subjective perception that can be modeled depending on life pace and life style. The Affective Health application should support reflection on life with e.g. a circular perspective that allows the user to look at her/his life separated from the standard units. One example of cyclic time is the ecclesiastical year that is not a year of moving forward, it is rather an eternal repetition of certain qualities of time; expectation, sorrow and sacred hope. Users should be able to scroll between reoccurring behaviors, events, emotions and so on. Only by exploring radically different alternative perspectives, such as circular time or time as relevant only vis-à-vis certain geographical positions, have we been able to move outside the prevailing perspective on time as linearly progressive in equal portions of time. [fig. 3] We also discuss how the time could be presented differentiated: People's lives are continually changing in tempo, intensity and activities. How can "important" events get more focus or space than "unimportant" and how would the user tell the application this?

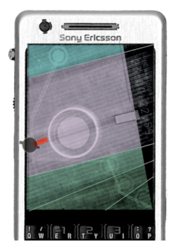

Fig 4. Sketch of a red flag to reinforce an event.

What we also want to dig deeper into is "primitive" time, point time, or moments of time where the future is impossible to imagine and where time is created from the pulse and life of humans. Memories or events that occur are tied to geographical grounds and not to a time frame.

Thoughts and further work

While we could have explored users' perception of time vis-à-vis stress through various user studies, understanding of cognition or through some theoretical models of subjective time, our approach has been to explore the time dimension through designing different representation. This helped us to clarify what it means to create a system where users can compare their subjective experience of stress over time, finding patterns in their own behaviors. We believe temporality is a core subject in our society that form people's lives and force a certain kind of living and perception of the world and ourselves embodied in this world. We who are born and raised in this structure are very much products of this and it is exciting to challenge this concept, and design for a broader perspective of time and temporality.

Acknowledgements

A lot of thank you to; Jarmo Laaksolahti, Pedro Ferreira, Claus Weymann, Christer Wigerfelt and Maria Lindström.

References

[1] Pedro Sanches. Supporting self-reflection in everyday life. Unpublished Master thesis.

[2] Gtbrg Stress Institute ref, will be completed.

[3] http://www.emwave.com/about_emwave_stress.html http://www.engadget.com/2005/11/10/nipro-cocoro-meter-the-saliva-powered-stress-detector/.

[4] Pedro Ferreira. "License to chill" How to empower users to cope with stress.

[5] Anna Ståhl. A foundation for Emotional Expressivity. Proc of 2005 conference on Designing for User experience. San Francisco, California

Discovering Temporal Categorical Patterns Across Multiple Records

Taowei David Wang[*%]
tw7@cs.umd.edu

Catherine Plaisant[%]
plaisant@author.ac.uk

Krist Wongsuphasawat[*%]
kristw@cs.umd.edu

Ben Shneiderman[*%]
ben@cs.umd.edu

*Human-Computer Interaction Lab
%Dept of Computer Science
University of Maryland
College Park MD 20742

CHI 2009, April 4 – 9, 2009, Boston, MA, USA
ACM 978-1-60558-246-7/09/04.

Abstract
This is a position paper for the CHI 2009 Temporal Workshop, summarizing our projects designed to specify queries across multiple records and provide comprehensible result presentation, especially for databases of Electronic Health Records.

Keywords
Categorical data, patient records, search, alignment, similarity

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

Introduction
An increasing number of temporal categorical databases are being assembled. For example millions of Electronic Health Records are being recorded by healthcare organizations, logs of traffic accidents are kept by transportation agencies, student records are maintained by academic institutes and the activities of suspicious individuals are being kept by law enforcement agencies. We are working on several projects which aim at improving the use of those databases by researchers and analysts who want to specific complex queries, explore patterns, make

hypotheses about recurrent sequences of events or find similar records.

Lifelines2:
Interactive exploration of multiple records Align, Rank, Filter and Group

Our previous work on Lifelines has shown that a timeline visualization for personal histories can provide benefits over a tabular view [1, 3], but many tasks involve temporal comparisons across multiple records relative to important events called sentinel events (e.g. a first heart attack). We explored strategies in supporting query specification, efficient search and comprehensible visualization in our Lifelines2 project [5].

Interactive visual exploration can complement query formulation by providing operations to align, rank and filter the results. Display of patient histories aligned on sentinel events enables users to spot precursor, co-occurring, and after-effect events. A controlled study demonstrated the benefits of providing alignment (with a 61% speed improvement for complex tasks). A qualitative study and interviews with medical professionals demonstrated that the interface can be learned quickly and seems to address their needs.

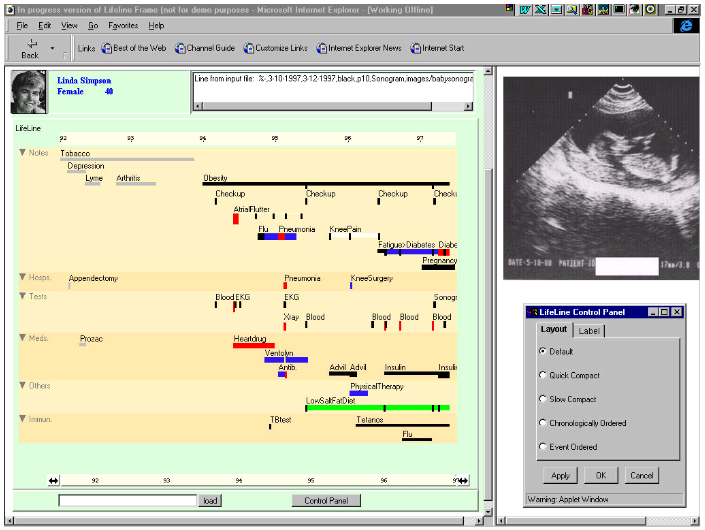

Figure 1. LifeLines present a summary of personal records—here, a medical record—on a zoomable timeline. LifeLines show multiple facets of the record, such as doctors' notes, hospitalizations, or tests, with line thickness and color used to map data attributes such as severity or drug dosage. For more information see http://www.cs.umd.edu/hcil/lifelines

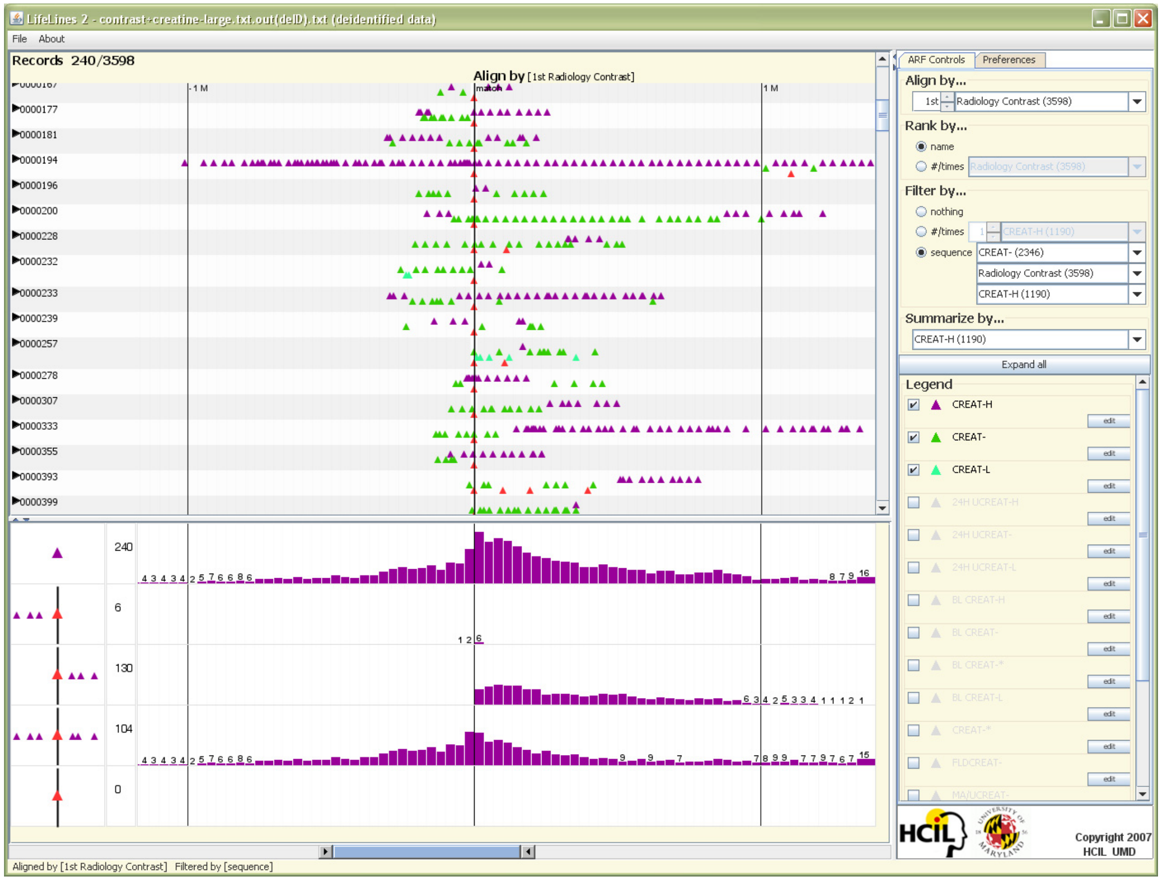

Figure 2. Lifelines2 allows the exploration of multiple records, which can be aligned, ranked and filtered. Temporal summaries allow comparisons of aggregate data across groups of records. For more information see: http://www.cs.umd.edu/hcil/lifelines2/

Pattern Finder

Current query tools make complex temporal queries difficult to pose, and physicians have to rely on computer professionals to specify the queries for them.

PatternFinder [2, 4] demonstrates a novel query tool implemented in a large operational system at the Washington Hospital Center (Microsoft Amalga, formerly known as Azyxxi).

Figure 3: PatternFinder. The user has specified the following query: "Find patients who had a normal serum creatinine lab test (the baseline event) less than 2 days before a radiology test with intravenous contrast (the sentinel event), followed by an increase in serum creatinine by more than 50% and of more than 1.0 mg/dl relative to the baseline measurement, within 5 days after the contrast administration". The query returned 12 patients. For each patient a timeline show the timing of events that match the query, aligned by the sentinel event. Each event is drawn as a color coded tick marks. In this query all baseline events appear on the left side of the sentinel event, and follow on are on the right side. Options allow users to display the numerical values as well. Zooming reveals more details. For more information see: http://www.cs.umd.edu/hcil/patternFinderInAmalga and http://www.cs.umd.edu/hcil/patternfinder

Similan: Finding similar records

It can often be helpful to find records that are similar to the record of a particular patient of interest. A major challenge is how to define a similarity measure that captures the searcher's intent. Many methods for computing a similarity measure between time series have been proposed, but temporal categorical records are different and require fresh thinking. In the Similan project we propose a temporal categorical similarity measure, called the M&M measure, which is based on the concept of aligning records by sentinel events, then matching events between two records. The M&M measure is calculated as a combination of the time differences between pairs of events and number of mismatches.

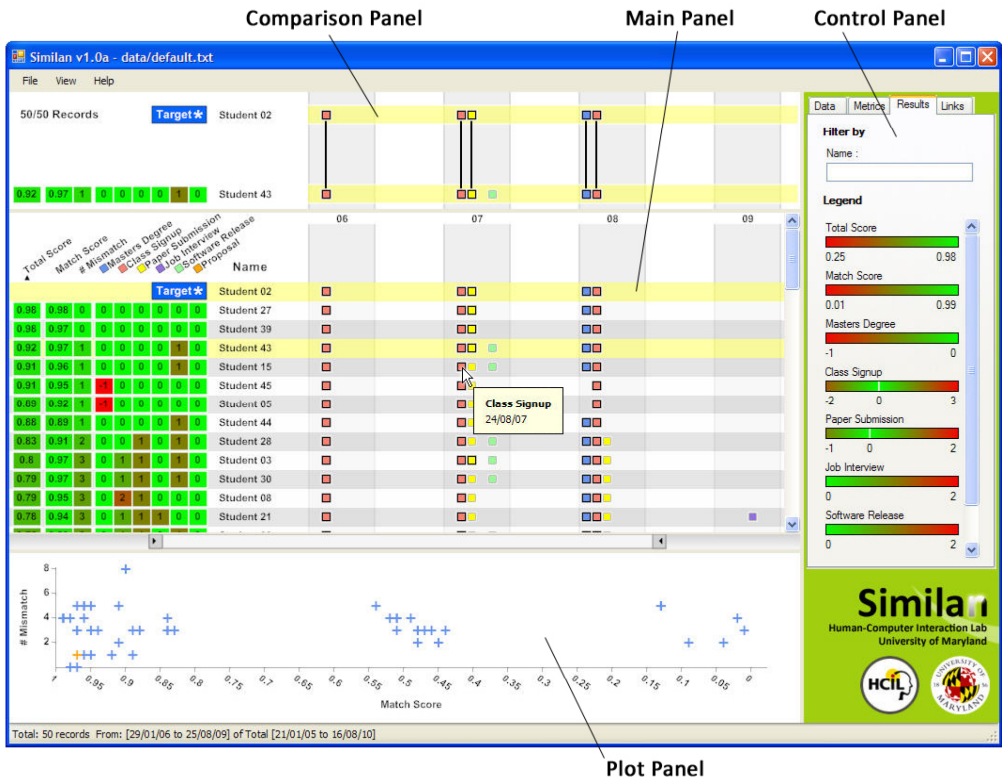

Figure 4: Similan interface allows users to specify a Target record and find close matches. For more information: http://www.cs.umd.edu/hcil/similan/

Acknowledgements

Our current work is supported by the Washington Hospital Center and focuses on Electronic Health Records.

Project Papers

[1] Alonso, D., Rose, A. and Plaisant, C., Viewing personal history records: A comparison of tabular format and graphical presentation using lifelines. *Behaviour and Information Technology*, 17, 5 (1998) 249–262.

[2] Fails, J., Karlson, A., Shahamat, L., Shneiderman, B. A Visual Interface for Multivariate Temporal Data: Finding Patterns of Events over Time, *Proc. of the IEEE Symposium on Visual Analytics Science and Technology* (2006) 167-174.

[3] Plaisant, C., Mushlin, R., Snyder, A., Li, J., Heller, D., Shneiderman, B., LifeLines: Using visualization to enhance navigation and analysis of patient records, *Proc. of the American Medical Informatics Association 1998 Annual Fall Symposium*. AMIA, Bethesda MD, 1998, 76-80

[4] Plaisant, C., Lam, S., Shneiderman, B., Smith, M., Roseman, D., Marchand, G., Gillam, M., Feied, C., Handler, J., Rappaport, H. (2008) Searching Electronic Health Records for Temporal Patterns in Patient Histories: A Case Study with Microsoft Amalga, in *Proceedings of the American Medical Informatics Association 2008 Annual Fall Symposium.* AMIA, Bethesda MD.

[5] Wang, T., Plaisant, C., Quinn, A., Stanchak, R., Shneiderman, B., and Murphy, S., Aligning temporal data by sentinel events: Discovering patterns in electronic health records, *Proc. of CHI 2008*, ACM, New York, 457–466.

Exploring Multivariate Data Streams Using Windowing and Sampling Strategies

Zaixian Xie
Worcester Polytechnic Institute
100 Institute Road
Worcester, MA 01609 USA
xiezx@cs.wpi.edu

Matthew Ward
Worcester Polytechnic Institute
100 Institute Road
Worcester, MA 01609 USA
matt@cs.wpi.edu

Elke Rundensteiner
Worcester Polytechnic Institute
100 Institute Road
Worcester, MA 01609 USA
rundenst@cs.wpi.edu

CHI 2009, April 4 – 9, 2009, Boston, MA, USA
ACM 978-1-60558-246-7/09/04.

Abstract

The analysis of data streams has become quite important in recent years, and is being studied intensively in fields such as database management and data mining. However, to date few researchers in data and information visualization have investigated the visual analytics of streaming data. Although streaming data is similar to time-series data, its large-scale and unbounded characteristics make regular temporal data visualization techniques not effective.

In this paper, we propose a framework to visualize multivariate data streams via a combination of windowing and sampling strategies. In order to help users observe how data patterns change over time, we display not only the current sliding window but also abstractions of past data in which users are interested. Uniform sampling is applied within a single sliding window to help reduce visual clutter as well as preserve the data patterns. However, we allow different windows to have different sampling ratios to reflect how interested the user is in the contents. To achieve this functionality, we propose to use a DOI (degree of interest) function to represent users' interest for the data in a particular sliding window. In order to visually

convey the multi-correlations and trends at the same time, we use multiple views, the union of traditional multivariate visualizations and line charts.

Keywords
Data stream, multivariate data, sampling, windowing

ACM Classification Keywords
H5.2. Information Interfaces and Presentation: User Interfaces—– Graphical user interfaces

Introduction

Advances in hardware enable people to record data at rapid rates, e.g. kilobytes per second or even higher speeds. Many real world applications require data collection at such a high speed. Moreover, the newly acquired or generated data items often need to be processed immediately. In the areas of database and data mining, the term *data streams* has been introduced to refer to such data that keeps growing and needs to be processed on the fly. Researchers have developed a lot of techniques to manage, query and analyze data streams and make decisions in real-time. But to date there has been limited contributions from the visualization community.

Since windowing strategies are commonly employed in the management of data streams, a naive solution to visualizing such data would be to split the whole stream into contiguous sliding windows and send each through the visualization pipeline one by one. The obvious disadvantage of this technique is that only the current fragment of the data stream is displayed, and it is difficult to learn how data patterns have changed over time because the past data is no longer visible. Another solution is using traditional time-series data visualization techniques. However, the commonly used abstraction techniques for static datasets, such as sampling, must be adapted to unbounded streams. Otherwise, the system will become less efficient if we keep restarting the data abstraction algorithm when new data items are available.

In this paper, we combine windowing and sampling strategies to preprocess multivariate data streams, and then visualize it using multiple views. We display not only the data in the current sliding window, but also abstractions of past data in the same or separate views, helping users study how data patterns change over time. For each sliding window, we allow different sampling ratios based on the degree of users' interest, which is determined by a DOI function (Degree of Interest) [1,2] defined by users.

We will focus on a special type of data stream, namely univariate-aggregation, in which each data item is a multi-dimensional vector and each dimension can be regarded as univariate data. This type is very common. One example is a stock price data stream in which each dimension represents the price of one company’s stock. Another is a set of sensors monitoring the status and positions of military vehicles. The main contributions of this paper include:

- We propose a framework to introduce windowing and sampling strategies into traditional multivariate and time-series data visualization techniques. This combination is aimed at handling unbounded input.
- This framework allows users to define a DOI function to describe the degree of users' interest for different portions of the data.

- Users can choose multiple views to observe the data stream, including traditional multivariate and time-series visualizations. Linked interactions among these views are provided to help users detect and isolate data patterns and their changes.

The Framework Based on Sampling and DOI Functions

Figure 1 shows the proposed framework. Here we use non-overlapped sliding windows. Because a data stream is infinite in nature, the data is sampled first and then saved in the *Data Item Memory Pool*. When the memory pool is full, a part of the old data will be moved to the disk pool. The *Mixer* is the core of the whole framework. Its input includes the current window and segments of old windows. It can assign sampling ratios to these windows and generate outputs that mix datapoints from these windows. Users can observe data from multiple views, such as scatterplot matrices for the representation of multi-correlations, and line charts to convey the trend of each dimension.

To determine the sampling ratios for current and old windows in the *Mixer*, we introduce a DOI (Degree of Interest) function to represent how interested the user is in seeing a particular sliding window. The DOI function has two parameters, a timestamp t_d for a specific sliding window, and the current time point t_c. Formally, the DOI function is given as: $DOI = f_{doi}(t_d, t_c)$. The DOI value will be mapped to a sampling ratio based on the rule that the portions in which users have higher interest will be shown in more detail.

In order to help users define DOI functions, we provide an interactive tool. In this tool, a curve corresponding to the DOI values for the past sliding windows will be displayed in an accompanying window. When the user is monitoring data streams, he can adjust this curve for more or less detail in specified portions.

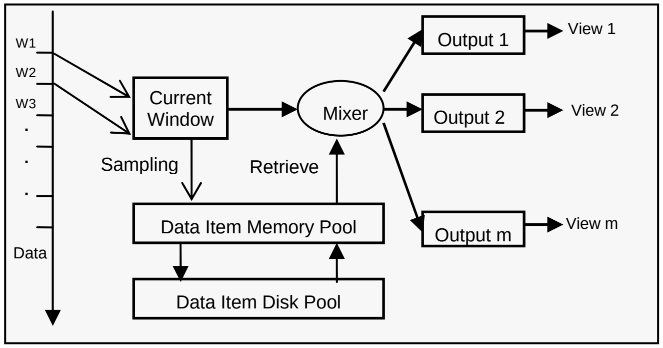

Figure 1. The framework of user-driven multiple-view visualization for data streams.

For convenience, we also provide two basic types of DOI functions for some common analysis tasks. One is used to observe how data patterns change across several sliding windows, namely the **RC** (Recent Change) Function. The other focuses on repeated phenomena, which we call **PP** (Periodic Phenomena) function. Figures 2 and 3 show examples of these functions.

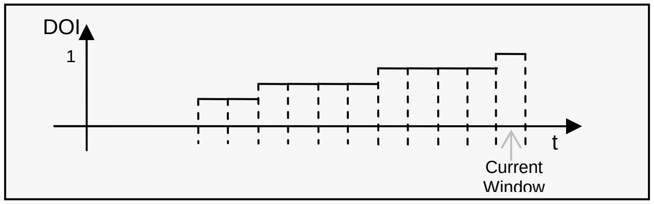

Figure 2. The DOI function for recent change.

Views

We propose three types of views for multi-resolution data streams after windowing and sampling.

- A general adaptation of traditional multivariate visualizations: We regard the output of the *Mixer* as a dataset and display it using traditional multivariate visualization techniques, such as parallel coordinates and scatterplot matrices. Available visual attributes are used to convey the timestamps of datapoints. For example, we can use dot size, color or opacity to show the age of data items in scatterplot matrices.

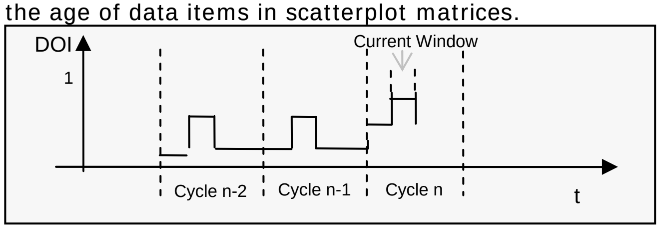

Figure 3. The DOI function for periodical phenomena.

- Embedded Views: We show line charts for each dimension and put them in the diagonal plots of a scatterplot matrix. This makes it easy to convey the trends for each dimension without the need for more canvas space (See Figure 4).

- Integrated Views: The disadvantage of embedded view is that time axes of all line charts are not in the same horizontal position, resulting in difficulties in comparing trends for different dimensions. We develop integrated views to display multiple views in the same canvas. Some views, e.g., line charts, are used to convey trends, while others could represent multi-correlations via multivariate visualizations.

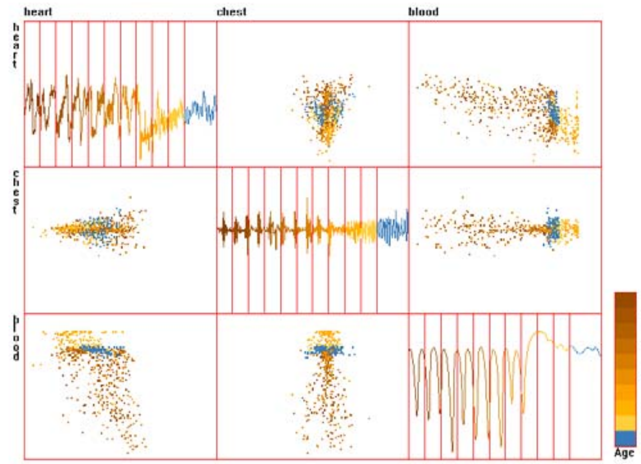

Figure 4 The dot color denotes the age of the datapoint (the difference between its timestamp and current time).

References

[1] Furnas, G.: Generalized fisheye views. Proc. ACM SIGCHI Conference on Human Factors in Computing Systems (1986), 16–23.

[2] Hao, M. C., Dayal, U., Keim, D. A., and Schreck, T. Multi-resolution techniques for visual exploration of large time-series data. EuroVis07: Joint Eurographics – IEEE VGTC Symp. on Visualization, pages 27–34, 2007.

Burn My Digital Remains: Life-Logs and Death

Jason Zalinger
Rensselaer Polytechnic Institute
110 8th Street
Troy, NY 12180
zalinj@rpi.edu

CHI 2009, April 4 – 9, 2009, Boston, MA, USA
ACM 978-1-60558-246-7/09/04.

Abstract
We need to sit down with our spouses and parents and children and talk about what we want done with our digital remains. The reason for this is because temporal data has reached a critical moment. Ubiquitous, cheap computing now allows a person to record every moment of their lives. One form of temporal data, "life-logging," poses a serious challenge for both the user and the user's loved ones. In short, what will we do with a lifetime’s worth of life data? Moreover, when you die, what will your family and friends do with your lifetime of audio, visual and textual data? This paper discusses some of the problems surrounding life-logs and death and offers some possible solutions.

Keywords
Temporal data, life-log, memory

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

Introduction
As is well known in the HCI community, in 1945, Vannevar Bush, wrote an article for *The Atlantic Monthly* called "As We May Think." Bush's main point is

that we have accumulated a giant body of research, but we have not figured out how to organize knowledge so that it is immediately accessible. With the war over, Bush asked, "What are the scientists to do next?" Bush's answer to this question was remarkable. He essentially sketched out what we now call the personal computer, the internet, Google and Web 2.0 all wrapped up in one, imaginative concept he called the "Memex," a device to record one's life and act as an "enlarged intimate supplement to his memory" [2]. However, Bush never considered what would happen to all this data once you die. The Memex does not account for our digital remains.

Should I Record All MyLifeBits?

The Memex continues to inspire creative projects in the 21st century. One of the most fascinating is the Microsoft funded, MyLifeBits, which describes itself as "a lifetime store of everything. It is the fulfillment of Vannevar Bush's 1945 Memex vision including full-text search, text & audio annotations, and hyperlinks" [5]. In a 2007 article, the creators of the project, Gordon Bell and Jim Gemmell, lay out their vision as a "quest to digitally chronicle every aspect of a person's life" [1]. After six years, one of the designers has captured more than 300,000 records, taking up about 150 gigabytes of memory [1].

Is this the "fulfillment" of Bush's vision, as the creators suggest, or is this a misguided corruption? What are the possible emotional pitfalls of human-life-as-database? If we are not worried about the impact on ourselves then we should ask what our loved ones will do with all these digital memories once we are gone. In short, if the dream of the Memex is embodied in MyLifeBits, shouldn't someone ask whether or not this is a good thing?

Bell and Gemmell argue that, "Human memory can be maddeningly elusive," and that "digital memories allow one to vividly relive an event" [1]. The assumption behind this statement is that humans *want* to remember everything, that somehow letting memories slip away is a problem. Don't we need to let certain memories fade? Isn't the fading of memory part of the healing and learning process? Bush proposed Memex with the goal of *organizing* knowledge, but MyLifeBits does not yet offer anyway to help the user or their family make sense of their digital life. Thus, MyLifeBits challenges our idea of self knowledge. Will we be able to move forward emotionally when our multimedia past is so accessible? How will grieving family members deal with fifty years of email, text messages and phone conversations? The dream of the Memex, as manifested in MyLifeBits and supported by a confluence of other technologies (Google, cheap storage, light and powerful digital camcorders, wide access to broadband) is not just changing our relationship to our present, but it has the potential to dramatically change our relationship with our past. It is our responsibility to take responsibility for our inventions. Computers have evolved from the distant, impersonal room-sized mystery boxes, to the personal computer and now we are on the brink of *intimate* computing.

The Problem With Life-Logs

In an extremely insightful article, Dodge and Kitchin argue that life-logs could provide "rich autobiographical narratives that are of potential use and value to the individual life-logee…" [3]. However, they also raise broader social and political questions. They argue that technologists have "strategically avoided" some of the "difficult ethical questions" concerning life-logs. In particular, they worry about who owns life-log data. Will third parties have access? Would those captured by life-logs have legal claims to access (e.g., family, friends, co-workers)? One of the most interesting questions they ask, however, is "what happens to the life-log at death? [3]. I am less concerned with issues of legal access, rather, I want to know how we can begin to help the owner of the life-log and their loved ones make sense of the "captured" past.

Possible Solutions

Dodge and Kitchin pose an interesting solution to the problems surrounding life-logs. They posit that forgetting is "not a weakness or a fallibility" but an "emancipatory process that will free life-logging from burdensome and pernicious disciplinary effects" [3]. They go on to propose that an "ethics of forgetting" be built into the design of life-log systems to ensure that humans still have the ability to forget the past because forgetting is part of moving forward emotionally [4]. I agree in spirit, but the implementation of an ethics of forgetting seems itself potentially unethical. By "tweaking" random life-log data to ensure a certain amount of "imperfection" we may do more harm than good. Rather, I propose another, far simpler solution—talking to each other.

First, we need technologists exploring life-logging to acknowledge the emotional, social, legal and ethical implications of their work. Second, whether you are logging every moment of your life or just a "normal" person producing staggering amounts of email each week, we need to sit down with our spouses and parents and children and talk about what we want done with our digital remains. When you die, should your wife and children have access to your passwords, or do you want your digital remains destroyed? Do you want it printed out? Do you want to "pass it down" to your children? Will certain friends be given certain chunks of data (road trips, parties etc.) Do I need to include a "digital clause" in my will? These are just some ways to begin talking about these issues. Whether it's a complete life-log of eighty years or ten years of email, we need to start including our digital history in conversations. My mother and father sometimes tell my sister and I which family artifacts we will get when they are gone. Grim thought it sounds, this is a common family ritual. We don't like to talk about death, but we have to. My argument is that, aside from the good china and mother's jewelry, we should also talk about dad's blackberry and mom's email archives. Rather than leaving my cremated remains in an urn, will a 20 terabyte hard drive sit on the mantle over the fireplace, wirelessly connected to a computer so that my family never has to forget *anything* about me?

Conclusion

Humans need to forget in order to move forward. By simply building life-log technology without any concern for the individual or their loved ones, we are making a terrible mistake and creating an impossible digital burden, and it will be left to our husbands and wives and children to figure out what to do with our digital remains. Life-logs could be a wonderful way to celebrate the life of those we love, but if we are not prepared, it might overwhelm us. It could become a powerful obstacle in the process of moving on. Time moves forward, but technology allows us to freeze it for a moment. But how many of life's moments are too many to leave behind?

References

[1] Bell, G., & Gemmell, J. (2007, February). A digital life. *Scientific American*. http://www.sciam.com/article.cfm?id=a-digital-life.

[2] Bush, V. (1945). As we may think. *The Atlantic Monthly*. http://www.theatlantic.com/doc/194507/bush.

[3] Dodge, M. & Kitchin, J. (2007). Outlines of a world coming into existence: Pervasive computing and the ethics of forgetting. *Environment and Planning B: Planning and Design*, *34*(3), 431-445.

[4] Lowenthal, D. (2001). Preface. In Forty, A. & Küchler, S. (Eds.), *The art of forgetting* (pp. xi-xiii). Oxford, UK: Berg Publishers.

[5] "MyLifeBits Project." http://research.microsoft.com/barc/MediaPresence/MyLifeBits.aspx.