Second Year Project Prospectus
First Draft
Gaston R. Cangiano
10/02/06
Cognitive Science Department
University of California at San Diego

Studying patterns of user organization in the personal computing space using ethnography and statistical modeling

Introduction and background

A very promising area of research in HCI is that of personal information management (PIM). A large segment of the research effort in this area is devoted to sophisticated methods of indexing and searching across hybrid types of data (e.g. email, text documents, web pages, and so on). Presently, personal computer users are required to organize their information according to a rigid structure imposed by the file system hierarchy, browser bookmarks, or any other application or operating system-defined structure. The state of the art around this challenge is to provide users with advance tools for searching or for organizing data according to some flexible paradigm. What we propose here is to address this problem by studying how people *actually* organize their information *while* they are using it. By doing so, we seek to uncover patterns of interactive behavior that can be used to programmatically assist users in organizing data. These behavioral patterns will be defined by the locus of activity (i.e., the current task), and applied in real-time, contrasting with present approaches (i.e. imposing the structure a priori). We envision a hybrid of searching and browsing through related information, called "contextual browsing", as opposed to scrolling though a list of search results. Therefore the main research question we will explore here is: what is the most appropriate *level* (i.e., window level, content level, mousing level) for capturing these "patterns" of user behavior? And how can we apply them in a useful fashion to aid organization.

The background motivation for this project comes from an area of research that started as early as the eighties, but that saw its most important implementation at the Palo Alto Research Center (PARC) in the nineties. The idea was that of replacing the present file system in personal computers with a fluid repository of data, where information could be organized semantically instead of a rigid structure. The "Placeless Documents" project in particular [Dou99], was the main seed for the new paradigm that motivates the research proposed here. We believe that a new paradigm such as the fluid organization of personal data requires us to take a new look at how information is used in context. We also believe that this new approach, in conjunction with the increasing sophistication of search technologies, can yield tremendous

results. This new approach requires new methodologies for research in Human-Computer Interaction (HCI).

There are several major software companies and research institutions applying a similar approach for slightly different problems. Microsoft, Intel and the University of Washington among others [DDJ+05] [GS00] [OST+06] [PFP+04] [PLF+03] have projects underway using statistical modeling to predict user activity and context. Our research proposal is similar in that it will employ statistical learning to recognize patterns of interactive behavior to assist users in the organization of their personal information. We will do so by collecting user activity histories at the window manager level. The main difference is that we will employ detailed observation (ethnography) of real-world activity to select the features of our statistical model in the first place. We believe this to be a significant contribution and novel approach.

Method

This research can be viewed as an ethnographic study of personal computer users in their daily activities. We are interested exclusively in users of windowing desktop systems (specifically Microsoft Windows, for implementation reasons to be described shortly), but in particular in their use of windows, task managers, application top level-content and shortcuts. In its first phase, this research will consist of video captures, interviews and direct observation. For the second phase, we are implementing a software data collection tool to collect digital data. Capturing desktop activity digitally will not only minimize the "observer" effect, but it will also provide a very rich data set to be employed in the later stages for statistical analysis.

We expect to collect data from at least 10-15 users. These users will be of varied backgrounds, and hence the nature of their recorded activity will vary greatly. Contrary to controlled laboratory methods, we want to capture users in their natural setting. Therefore we will *not* run our data collection in a laboratory setting with a given predefined task, but rather we will go to locations to observe and collect data from users everyday activities. This factor adds to the challenge of our problem but it also increases its validity and theoretical generalization potential from the perspective of our laboratory (Distributed Cognition and HCI at UCSD). We have started our data collection on several participants (4 so far). A video recorder was setup for a period of two hours, recording only the activities on the computer screen. For the second phase of our research, we expect to install our data collection tool and run it for a period of a week per user. Users will have the ability to easily disable the tool for privacy concerns. The digital collection tool will also record computer screen visual captures at a regular frequency so that the digital data can be visually correlated and verified by the PI.

The specific and final nature of the digital data to be collected will depend directly on the implementation feasibility of the data collection tool. It is currently being implemented in Microsoft's Windows XP operating system, and we expect to be able to collect data on mouse, keyboard and windowing operations. The Windows operating system provides access to a list of the current windows on the desktop, as well as messaging events between the OS and those windows. Furthermore, certain information is available from these window handles, such as application type, document title and location. We are interested in capturing data at the window manager level, including semantic content of the user activity (i.e. document top-level content, such as URLs , titles and metadata). The details are yet to be determined with further research into this implementation.

It is important to note that the first phase of this research project is directly complementary to the second phase. More specifically, we seek to employ our ethnographic observation for selecting the most significant variables (i.e. features) for our statistical analysis, and later on, statistical modeling of user behavior. We strongly believe this combination of methods is a unique and much needed approach in HCI. Most current research of this ilk, hinges on classic statistical methods for selecting the features for the models. Researchers usually start with their own knowledge of computer usage (e.g. the task manager in Windows OS) and then use standard techniques in statistical modeling for reducing uncertainty and selecting the most appropriate predictive variables. In contrast, what we propose here is to employ direct observation to select the features initially, and only at later stages employ statistical methods of space reduction. We strongly believe this unique approach will yield significant results in what is often called the "art" of feature selection in statistical modeling, often relying on expert advise.

Expected outcome and applications

The projected outcome of this research is to first and foremost collect a significant amount of rich data, mainly digitally, about *activity histories*. We expect to learn about several styles of interactive behavior varying considerably across users and across tasks. We are confident that we will find commonalities across these styles. Perhaps these commonalities will not be perceptible to the naked eye and require fine statistical analysis. Nonetheless, the observation portion of the research will reveal to the PI the potential commonalities. We will perform interviews with our participants, asking questions that can be used to validate the results obtained from correlating the digital data. For example, we anticipate finding very useful information in the manipulation of overlapping windows. For instance, a user could partially overlap windows in order to cut and paste between them, or for the purposes of having simultaneous exposure to the content of those windows. In a case like this, we will study whether there exists significant correlation between this windowing pattern and the *content* of those windows. We will analyze the correlation statistically and validate it against the responses from the user. Similar results can be obtained

from looking at window toggling and mouse movement patterns, and the type of the content being manipulated or edited (IM, email, doc, etc).

In the ideal scenario, a user would perform a combination of *searching and browsing*, instead of searching exclusively. Researches have already begun exploring this idea under the rubric of "pivoting" [GLB03]. We envision this new paradigm as "Contextual Browsing", a hybrid of searching and browsing through the context of activity instead of a list of results. A user starts with either a search query or by pulling up a document; he can then navigate through a web of related documents and data bits narrowing the query, or simply move into a different context, say a different task or different date, to then continue navigating from that new point (or pivot). We are currently working with a group of Computer Science master students at UCSD, who are in the process of creating the infrastructure and the front end design (GUI) required to implement a related organization paradigm. We expect to collaborate with them throughout the near future in order to enhance our research perspective, validate and test our results.

In the long term, our research will produce a robust statistical engine that will run in the background of an operating system's window manager. This engine will consists of a predictive model that learns from user activity and associates related data as well as automatically adds tagging data (metadata). The model will be designed based on the results from the proposed research presented here. One of the main contributions of this engine could be acting as a "suggestion engine", allowing users to link related data, and assisting them in the contextual real-time organization of their personal information. We plan to validate the design of this model by conducting studies with human participants. The validation tests can be performed in a laboratory setting, having participants perform a task while running the predictive model in the background. The predictions made by the model will then be tested against the users knowledge via a questionnaire. A robust accuracy measure can be obtained in this fashion and the model can be tuned and re-tested iteratively.

References

[DDJ+05] Dragunov A., Dietterich T., Johnsrude K., McLaughlin M., Li L. & Herlocker J. 2005. *TaskTracer: a desktop environment to support multi-tasking knowledge workers.* Proceedings of the 10th international conference on Intelligent user interfaces. San Diego, California.

[Dou99] Dourish P. 1999. *Presto: an experimental architecture for fluid interactive document spaces.* ACM Transactions on Computer-Human Interaction, 6(2).

[GLB03] Gemmell J. ,Lueder R. & Bell G. 2003. *The MyLifeBits Store*. ETP. Berkeley California.

[GS00] J. Goecks and J. Shavlik. Learning users' interests by unobtrusively observing their normal behavior. *Proceedings of the 5th international Conference on intelligent User interfaces.* New Orleans, Louisiana, United States, January 09 - 12, 2000. ACM Press, New York, NY, 129-132.

[OST+06] Oliver N., Smith G., Thakkar C. & Surendran A. 2006. *SWISH: semantic analysis of window titles and switching history.* Proceedings of the 11th international conference on Intelligent user interfaces. Sydney, Australia.

[PFP+04] M. Philipose, K. P. Fishkin, M. Perkowitz, D. J. Patterson, D. Fox, H. Kautz, D. Hahnel, Inferring Activities from Interactions with Objects, *IEEE Pervasive Computing*, Vol. 03, no. 4, pp. 50-57, Oct-Dec, 2004.

[PLF+03] D. J. Patterson, L. Liao, D. Fox and H. Kautznferring. High-Level Behavior from Low-Level Sensors. *Proceedings of the Fifth International Conference on Ubiquitous Computing.* Seattle WA October 2003. Springer-Verlag Berlin, Germany.