

AppInsight: What have I been doing?

Gabriel Barata
INESC-ID/IST
Rua Alves Redol, 9
1000-029 Lisboa, Portugal
+351 21 3100289

gabriel.barata@ist.utl.pt

Hugo Nicolau
INESC-ID/IST
Rua Alves Redol, 9
1000-029 Lisboa, Portugal
+351 21 4233567

hman@vimmi.inesc-id.pt

Daniel Gonçalves
INESC-ID/IST
Rua Alves Redol, 9
1000-029 Lisboa, Portugal
+351 21 3100248

daniel.goncalves@inesc-id.pt

ABSTRACT

In this paper we present *AppInsight*, a visualization tool that enables users to reminisce on their computer usage history and derive meaningful insights about behaviors and trends. Human memory has the ability to re-experience episodes from our lives when supplied with suitable contextual cues, such as places, music, and so on. We explore a small set of properties, such as the application's name, URL and window title as contextual cues, in order to characterize the users' activity on their personal computers and how it evolves over time. Our user study shows that users enjoyed viewing their computer usage history and were able to both recall past events and introspect about their lives. Moreover, one of the most surprising outcomes is that they found several different applications for our tool, such as improving usage behaviors, controlling productivity, generating activity reports, and monitoring users in psychological studies. Finally, we discuss some lessons learned from our study and propose future research directions.

Categories and Subject Descriptors

H5.2. Information interfaces and presentation: User Interfaces.

General Terms

Design, Experimentation.

Keywords

Information visualization, desktop activity, computer usage.

1. INTRODUCTION

Nowadays many of us spend most of our time in front of a computer, be it for working, leisure or to communicate with others. As a matter of fact, we might very well spend more daily time interacting with such digital devices than with any other tool or person. Consequently, we began to generate more and more personal information which soon grew wild and became almost impossible to deal with. To make matters worse, digital storage also became cheaper and “stuff” that goes into the computer does not come out anymore [5].

The retrieval of personal information, such as email, photos, music, videos or spreadsheets, has proven to be a daunting task and, therefore, came to be an important topic of research. Several

forays have been carried out to study how to annotate documents and inter-relate them with other objects and people [7, 9]. The patterns and trends within these collections of personal bits have also been widely explored, especially for email [20], music [2] and web activity [1], as they convey much of our personal experience and, therefore, can be easily explored to promote reminiscence of the past and introspection.

These systems aim to be a memory aid by exploring our episodic memory [17], which allows us to re-experience specific episodes from our past and thus make sense of what happened and better get to know ourselves. A key component to help us recall those episodes are what may be called of contextual cues, which may be activities, places, events and other things related to what we want to remember [4]. Wherefore, contextual cues such as a person we have met or a visited website can be used to retrieve a document that was handed to us during a meeting [11].

Although the visualization and retrieval of personal information has been a major concern, little attention has been paid to the daily usage we make of our personal computers, which can tell so much about us. Few works went beyond studying low-level user interactions, such as window switching.

In this paper we investigate how the visualization of one's computer activity may impact reminiscence of the past on a specific time period and thus potentiate introspection. Due to the huge amount of data and the difficulty of processing it in real time, we also propose a limited set of contextual cues that can easily be extracted and are still able to convey usage of applications and websites over time. In order to carry out this study, we deployed an unobtrusive tool, named *AppAware*, to silently record every used application and visited website throughout three months. In the end of this period, a visualization tool, named *AppInsight*, was supplied to the users in order to allow them to make sense of their data. We present a few use cases to better illustrate how the visualization affected the users.

2. RELATED WORK

Much research has been done in information visualization in order to convey users with means to assess their data. In the following subsections we describe relevant works that aim to portray patterns and trends about several aspects of the user's history.

2.1 Document Visualization

In order to explore rich trends and patterns from the users' personal information, several works focused on specific types of documents. PostHistory [18] and Themail [20] aim to portray relationships by analyzing interactions between people in email archives. MUSE [12] has gone a little further by exploring additional contextual cues and automatically inferring social groups. Similarly, YouPivot [11] also took advantage of rich contextual cues to allow users to “*search for contextually related*

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

AVI '12, May 21-25, 2012, Capri Island, Italy

Copyright © 2012 ACM 978-1-4503-1287-5/12/05... \$10.00.

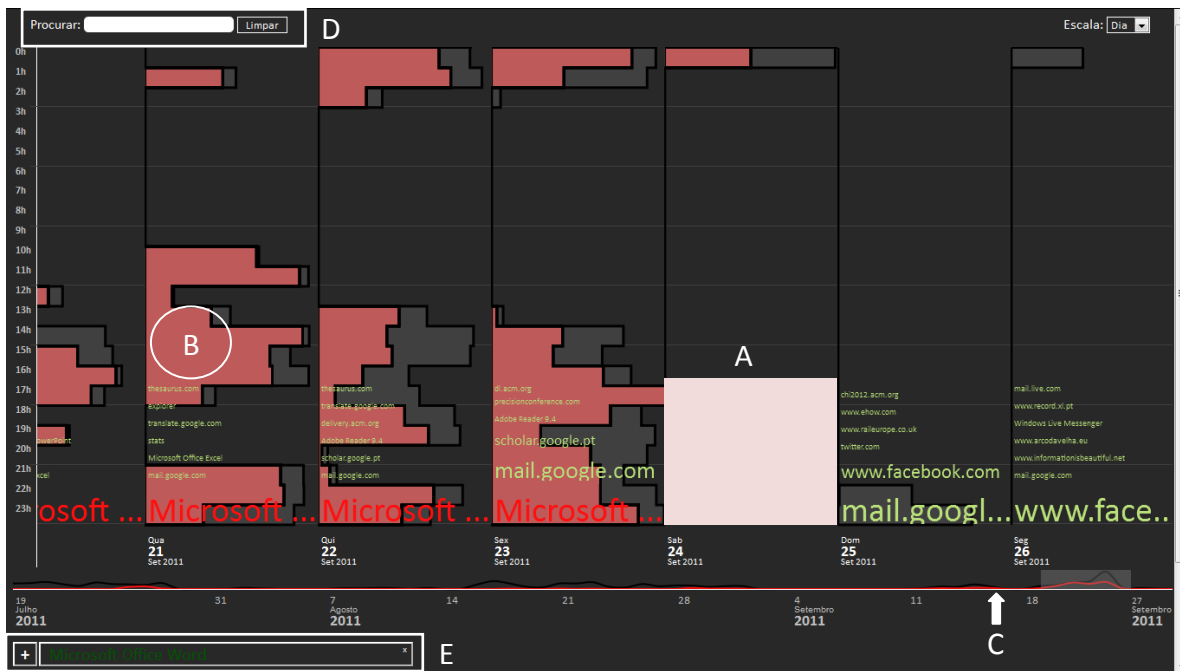


Figure 1. AppInsight is composed as follows: a) each column displays the most used applications per day and b) the overall and application specific time allocation per hour; c) a timeline where both the overall and application usage during the whole period are shown; d) a search box, and e) a zone that displays a pane with overall stats for each selected application.

activities and find a target piece of information". On the music scope, Torrens and Arcos [16] propose a few techniques to facilitate exploring music playlists and discovering the listening patterns therein. LastHistory [2], on the other hand, displays music listening histories over time, which allows users to analyze listening trends and to recall the past. Graham et al. [10] explored chronological display and summarization of photos, by taking advantage of photo metadata and clustering algorithms. Time-Quilt [13], in turn, allowed users to explore their personal photos by using a zoomable timeline. All of these approaches depict the users' histories regarding specific types of document usage and offer, thereby, support for introspection and retrospection. Nonetheless, they do not address, nor were they intended to, computer or web activity history.

2.2 Desktop and Web Activity

Currently, there are several commercial services available that measure time spent by users on each application and allow them to define activity goals¹ or to have their productivity monitored by their superiors². Nonetheless, none of these solutions provides a visualization of the overall usage history that can facilitate introspection.

This was tackled by works such as the one of Beauvisage [3], who investigated computer usage of a large panel of French users and was able to determine that they developed highly targeted application usages. Nevertheless, this work focused on general population trends rather than personal usage patterns. DejaView [14], on the other hand, records users' desktop computing experience and allows them to browse, search and revive personal records. Even though they can time travel through their computing activity, no further statistical representation is provided, making it harder for them to take a glance on the overall picture of their

history. Probably the most similar research effort to our own work is Window Watcher [15], which logs every window switch made by the user and provides a visualization tool to represent application usage and window placement over time. However, Window Watcher was developed to help researchers understand how users manage their windows, and was not designed to promote any kind of self-understanding.

Activity on the web has been also explored on works such as PeopleGarden [23], which uses the flower metaphor to depict one's posting history and the overall social environment on message boards. Posting history in newsgroups [19] and Wikipedia [21] have also been addressed in other works. Web browsing activity is also a topic of great interest to web designers which have studied how people navigate within websites and which paths they take [8, 6, 22]. Aquin et al. [1], on the other hand, presented a set of tools to monitor and visualize web activity from the personal standpoint, thus allowing users to grasp a deeper insight on their own online behavior.

Although presented works provide innovative techniques to depict the users' digital life, none covers their personal activity in the computer holistically (including both native and web applications) without requiring user intervention. Our goal with this work is to empower users with means to grasp insights on their behavior and reminisce the past. We developed *AppAware* and *AppInsight* to bridge this gap.

3. APPINSIGHT

AppInsight is a tool to visualize and explore one's activities on personal computers. The interface was designed to allow users to get a sense of their overall usage patterns over time, thus enabling them to reminisce past events and promote introspection behaviors.

Our goal was to help computer users to answer three main questions about their own activity: *What sort of things do I do on my computer? When do I do it? How all this evolves over time?*

¹ <http://www.slifeweb.com/>

² <http://www.softactivity.com/employee-monitoring.asp>

AppInsight has three main linked visualization areas arranged along a timeline (see Figure 1 a, b and c): *applications* – columns of words – along the timeline; *hourly usage* – hour chart – displayed in the background; and *usage evolution* – line chart – that appears at the bottom.

Applications reveal the most used applications in a particular period of time. Since we consider both native and web applications, they are displayed by their name or URL accordingly. For simplicity, from now on, applications refer to both native and web based applications. Their names are displayed in columns of words each with its own size. They are listed in a descending order, starting at the bottom. Most used applications are placed near the timeline in order to provide a more efficient visualization when users are exploring their data.

The selection and font size of an application’s name is based on the time users spent using that particular application. Font size varies linearly from 10pt to 40 pt. For instance, if application “xpto” is used most of the time, the name will appear fairly large when users visualize their usage history. The remaining applications will appear relatively small in that time period. If, on the other hand, many applications were equally used, they will all appear with similar sizes, exhibiting an idea of equal importance. Our goal is to present users with a meaningful visualization; one that they can relate to and easily infer their activities.

Hourly Usage is displayed in the background and exposes whenever there was some activity in the computer. As with most calendar interfaces, an hour-based view is exhibited where each bar corresponds to the amount of activity in that time period. The more activity there is in an hour, the bigger it gets. This visualization area enables users to get an overall view of their day and identify the most productive time periods. For instance, if the computer is an essential tool at work, it is expected a regular and consistent activity during the person’s work hours. If, on the other hand, the computer is used only at home for leisure purposes, after a long day of work, possibly a more sporadic and less consistent pattern may arise at night. Because user activity is presented over time, it becomes very easy to spot periods of intense activity as well as days when there was no activity at all. In other words, overall usage behaviors (and exceptions) are easily identifiable.

Usage Evolution is displayed as a line chart, placed above the timeline, and represents the progress in users’ computer usage. This specific part of the interface intends to inform the user about their overall activity giving a sense of continuity. Moreover, its placement above the timeline enables the users to easily perceive and explore the most relevant time periods.

Overall, *AppInsight* was designed to reveal patterns of computer usage that illustrate different behaviors for its users. Each pattern can be based on one or several attributes, such as:

Assiduity: *Sporadic* – low number of days active, typically do not spend much time using the computer and when it does, it does not follow a specific pattern – vs. *Regular* user – moderate to high number of days active and uses the computer frequently.

Diversity of applications: *Focused* – low number of applications, usually indicating that the user perform repetitive tasks or is concentrated in one task – vs. *Diverse* – high number of distinct applications, possibly not following any routine

Time allocation: *Continuous* – computer usage is restricted to a well-defined period of the day – vs. *Fragmented* - usage is distributed along all day, showing that the user does not use the computer in an intense manner (e.g. due to other non-computer tasks or breaks).

Type of applications: “*Workaholic*” – computer is mostly used for work purposes – vs. “*Slacker*” – person uses the computer just to spend some time or for recreation purposes.

Although we explicitly identify some patterns that can be revealed by *AppInsight*, we also acknowledge that others or some sort of combination may arise, since computer activity is very personal and unique.

3.1 Interacting with *AppInsight*

AppInsight displays the computers’ overall usage over time. Users can switch between three types of scales: day, month, and year. This provides them the opportunity to visualize their data with different levels of details.

Overall activity is displayed on the background and the most used applications are listed within the columns in the foreground in order to provide some clues about the users’ activities. Information about the day (both day of the month and week), month, and year is displayed at the bottom of each column, as well as below the timeline. Users can drag the timeline bar in order to browse their usage history and get further details. *AppInsight* also provides a search box, which allows users to find specific applications. As users type the name they are searching for, application names are instantly highlighted in columns and small bars appear above the timeline indicating a match (see Figure 2). Users can drag the timeline’s bar or directly click the timeline to jump to the desired time period.

3.1.1 Getting details

In addition to visualize their overall patterns, users can get further details about their computer behavior by interacting with applications’ names.

Since some application names cannot fit in a column, due to their string size and font, their names are truncated and the token “...” is added to indicate that the name is not entirely displayed. Nevertheless, users can bring up more details about applications. Whenever a user hovers the mouse over an application name, several actions take place: 1) the full name is displayed as a small

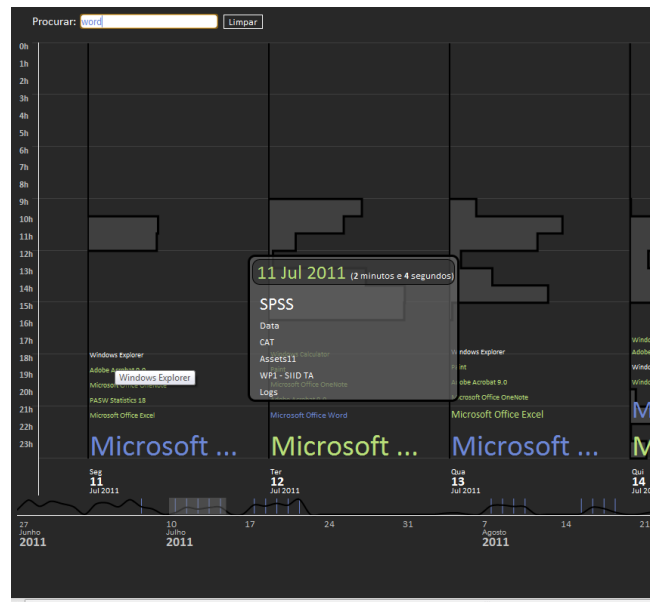


Figure 2. Searching for “word” on *AppInsight*. Notice how every day in which Word was used is highlighted on both the timeline and the day columns. Also, when the user hovers the mouse on an application, an information box is displayed.

label; 2) the name is shown highlighted in all visible columns, enabling users to see the progress of that particular application, even when it is displayed in undersized fonts; 3) an information box appears and displays the total time of usage for that day/month/year along with the most common window titles of that application (see Figure 2); this allows users to recall the context in which the application was used in that period of time. For instance, if a user hovers over “Microsoft Word”, the most used document files would appear, ordered by usage time. The font size is calculated by taking into account the time spent working in every document.

In addition to mouse hovering, users can also select applications. By clicking a name, the background chart highlights the usage of that application in all visible columns (see Figure 1 b). This allows users to get a more detailed visualization of their computer usage behaviors and how it is related with their applications. For instance, they can visualize at what time of the day each application is used the most.

When a selection occurs, the computer usage evolution chart, above the timeline, is complemented with the application’s usage; this enables users to relate the overall usage of the computer with a specific application. Also, some overall information and statistics about the selected application are placed in the bottom of the interface in a collapsible box (see Figure 1 e), featuring: full path, total usage time, mean usage time, weighted mean usage time, most used day/month/year, most used hour of the day, and most used titles.

3.1.2 Stacking & comparing applications

To provide a way to get further details about some activities, *AppInsight* allows the selection of several applications simultaneously. By pressing the “Ctrl” key, multiple selections can be made by clicking application names. To avoid overloading users with too much information, we restricted the number of simultaneous selections to three. During our experiments with early versions of *AppInsight*, we found this to be a good compromise between overload and required information.

Multiple item selection allows users to compare applications both within a specific period of time (e.g. day) or its overall usage. Thus, they can easily see the relationship between applications, e.g. complementary (used at the same time), interchangeable (one application is used in detriment of other) or dominance (one application is always more used than other).

3.2 Capturing and Processing the Users’ Activity

Monitoring users’ activity on a computer is a complex task and poses two main challenges:

1. Data collection must be efficient and should not overload the user’s machine. The number of possible applications to be monitored is manifold and this number grows even larger if we consider web applications.
2. Data collection must be non-intrusive, as users should not be disturbed or distracted at any time.

In order to overcome these obstacles, we developed *AppAware*, a non-intrusive lifelogging tool for Microsoft Windows, using C#, which continuously monitors interactions between the user and any application. To surmount the second challenge, *AppAware* sits in the background the whole time, watching users interacting with applications, without showing any notifications, pop-up messages or other manifestations that might distract them.

Reducing the computational load of watching every user step is a requirement for such an imperceptible watcher, and thus we decided to collect data only when the window taking focus changes or when it is refreshed, i.e., a tab is switched. This allows us to pay attention to a small set of operating system events and to have an application independent approach. We also had to go deeper and discover a reduced set of contextual cues, that could be gathered every time the focused window changed and would still enable users to reminisce the past and retrospect about their lives. We opted for the application’s identifier and the application’s window title, since these are two attributes that generally characterize the users’ activity.

The Application’s identifier may be derived from two distinct sources: in case it is a web application, the domain URL; otherwise, the executable file description is used. For example, if the user is browsing Facebook, the application identifier is “www.facebook.com”. On the other hand, if the user is using Microsoft Office Word to write an essay, the identifier is “Microsoft Word”. The application’s identifier is the most relevant memory cue, as applications are the basic building blocks of one’s computer activity history.

The Application’s window title offers users a summary of what they were doing with a specific application on a given time. For instance, in Facebook, the window title reveals whose page we were visiting or whether we were browsing our private messages. In other applications, such as Microsoft Word or Excel, the title displays the name of the documents we were working on. This provides users with a valuable insight about specific usages of certain applications, allowing them to identify documents, events and people.

Along with the application’s identifier and window title, the timestamps are also recorded every time a window switching or refresh takes place. This information is then processed and exported as usage statistics in a JSON file, which can be visualized in *AppInsight*. This visualization tool was developed as a web application, using Protovis³, in order to allow users to explore and make sense of their information using their web browser of election.

Three main statistical measures are computed to convey information about application usage history:

Application usage ratio (see Figure 1 a) is used to identify the most used applications for each day, month or year, which are depicted by each column in the visualization. Let *usage* be the function that returns the usage time for an application on a given periodic unit *p*, such as an hour, day, month or year, and *apps* an array of size *n* containing all applications the user interacted with. The usage ratio *r* for the *i*-index application is computed as follows:

$$(1) r(apps[i]) = \frac{usage_p(apps[i])}{MAX(usage_p(app[1]), \dots, usage_p(app[n]))}$$

We also compute application usage ratio for each hour, consequently displaying the resulting values as colored bars in the background (see Figure 1 b).

Overall application usage ratio represents the usage of an application over the course of the whole period, during which the user was observed. Let *totalusage* be the function that returns the total usage time of a period *p* and *m* be the number of periodic

³ <http://mbostock.github.com/protovis/>

units therein. The overall usage ratio o for the i -index application can be computed with the following formula:

$$(2) o(apps[i]) = \frac{\sum_{j=1}^m usage_j(apps[i])}{MAX(totalusage(1), \dots, totalusage(m))}$$
$$(3) totalusage(p) = \sum_{j=1}^n usage_p(apps[j])$$

This measure is represented by the colored line chart displayed right above the timeline (see Figure 1 c).

Overall application statistics are computed in order to provide users with a deeper insight over their interactions with a specific application. Every time the user selects an application from the interface, a panel appears below the timeline (see Figure 1 e), which upon clicked will unfold and display usage statistics. For the daily view, those are total usage time, average usage time per day, most used day, most used hour, most common titles and the application's path.

4. EVALUATING APPINSIGHT

Our goal with this work was to allow users to view their overall usage patterns over time. We wanted them to be able to explore their data, while reminiscing about daily activities and grasping insights about personal trends.

Since each user has its own unique usage history, the evaluation had to be made "in the wild". Unlike other types of data, such as email, where users have their archives of received and sent messages [20], there is no built-in tool on current operating systems that saves computer activity.

Rather than build a common dataset for all users, we decided to deploy our monitoring application (*AppAware*) on the users' personal computers. This meant that they were able to visualize their own application usage without having to be concerned about the privacy of their data.

4.1 Research Questions

This user study aims to answer several research questions regarding our visualization tool:

1. Is *AppInsight* useful?
2. Does *AppInsight* allow users to reminisce about past events?
3. Do participants enjoy using *AppInsight*?
4. Would users change their behavior because of *AppInsight*?

4.2 Procedure

This user study had two main stages: monitoring and visualizing. At the beginning of the first stage, participants were told that the overall purpose of our work was to provide a mean of analysis of computer-based activity. We then conducted a questionnaire to gather some demographic data and asked participants to install the monitoring application (*AppAware*). We explained that the application would monitor all used applications, including web browsing. We stressed out that all data would be saved exclusively on their own personal computers. No data would be sent by network or used without their explicit permission. Nevertheless, participants were still able to disable the monitoring process through our application menu.

In this stage, we decided not to provide the visualization application – *AppInsight* – in order to control previous experience with it. Since some participants could use it on a daily basis, while others might never visualize their data, this approach guaranteed that participants had the same experience level with *AppInsight* for the second stage of the user study.

After three months of monitoring we started the second stage: visualizing each participant's computer usage history. We started by installing a new version of *AppAware* that allowed participants to process and visualize their data using *AppInsight*. We then explained how they could explore their data with our visualization tool, and asked them to browse their usage history. Participants were encouraged to think aloud, share opinions, and describe interesting events. The interviewer interacted with participants through a semi-structured interview. After using *AppInsight*, participants answered a debriefing questionnaire about their experience with the tool. The entire procedure of the second stage lasted approximately one hour for each participant.

4.3 Participants

Recruiting users to get involved in this work proved to be an overwhelming task. Since we needed to gather the users' usage history, they highlighted many privacy concerns about participating in the study. Although we guaranteed that all data would be stored exclusively in their personal computers, for their own use, people still raised a lot of questions. In the light of these difficulties, we were still able to recruit 10 participants, all of them acquaintances of, at least, one of authors. The only requirement to participate in the user study was that they owned a personal computer and used it on a regular basis. Three of the participants were working professionals, while the remaining were students (five undergraduate and two postgraduate) from different universities. Of all participants selected for the user study, three were female and seven were male, with ages ranging from 19 to 30 years old. Seven of the participants used their computers for work purposes while three used for entertainment only activities (e.g. gaming, messaging, web browsing).

5. RESULTS

On three months of monitoring, participants had on average 15 days of uninterrupted computer activity, which corresponds to an average of approximately 4 hours and 30 minutes of usage per day. Nevertheless, computer usage per participant was diverse, ranging from 4 to more than 75 days of uninterrupted usage, illustrating very different usage behaviors. All participants were very keen to use our monitoring application, since only half of them ever disabled it. Even for those, the disabled time was on average 11 seconds, which is an insignificant value within the participants' activity history. This means that they were able to view their computer activity using a reliable dataset.

When asked about *AppInsight*, using a 7-point Likert scale, being 1, strongly disagree, and 7, strongly agree, the median [Interquartile Range] for the debriefing statements was:

"AppInsight is useful" – 5.5 [1.75] – Most participants found *AppInsight* useful, although some saw bigger advantages on using this kind of visualization than others. Particularly, participants that used their personal computers for work purposes found this tool very useful to monitor and evaluate their productivity – *"This is useful for time and productivity management, at least for me since I use my PC mostly for work. I think it is a great tool"*. Nevertheless, classifications were always positive.

"It was easy to understand the presented information" – 6.5 [1.0] – Overall, participants found it easy to use *AppInsight*. Activity was easily understandable using the applications' names and usage during the day. Participants effortlessly related to their activity history and immediately started reminiscing and interacting with the tool in order to find new information – *"the overall activity in the background ... it is fast and easy to see and gives me an immediate idea of what I did that day"*.

“It was fast to find the desired information” – 6.0 [1.0] – Most participants found that using a temporal arrangement of their activity allowed them to find the required information in a timely manner when searching for a specific day or month. The timeline enabled them to efficiently browse their data while skimming through their activity history. Also, when searching for a specific application, participants used the always available search box.

“AppInsight presented all information since it was installed on my computer” – 6.0 [3.25] – Although participants were generally impressed by the completeness of their computer usage history, some were also quick to point out some critical shortcomings. For instance, even though our monitoring application (i.e. *AppAware*) saves all applications’ names and titles, *AppInsight* was developed to show the overall pattern of users’ activity. Therefore, it only shows the most used applications in a period of time. We discuss this and other findings in more detail in the section entitled *Discussion and Lessons Learned*.

“I enjoyed seeing my computer usage history” – 7.0 [1.0] – Most participants (60%) strongly agreed with this statement and the remaining (40%) classifications were always positive. This shows that participants really enjoyed seeing their computer usage history using *AppInsight*. While some were more focused on reminiscing past events and activities, others were all about productivity. Yet, browsing computer history has shown to be an interesting task that most users enjoy. For instance, participant 5 said, *“I really enjoyed looking at the overall patterns. The daily activities are not that interesting for me, since I am very irregular, therefore it is not very useful, but with the monthly view one can see the overall patterns of activity and lifestyle”*.

“AppInsight allowed me to learn new information about myself” – 6.0 [3.25] – When using *AppInsight* participants easily remembered their past activities. Therefore, even though most classifications were positive, some participants stated that all shown information was something that they already knew – *“There is nothing here [AppInsight] that I find weird, and there is nothing missing. This corresponds to my mental model. Probably I would not be able to recall all these events, but once I saw them ... yes ... it makes sense and it does not surprise me”*. On the other hand, a number of participants were surprised about their computer usage – *“I can clearly see that I have a work routine and it is consistent! I am so happy ... wish by boss was here to see this. All leisure applications, such as facebook, are used after work or at my lunch break. Auto discipline is very important for me.”* In this case the participant really learned something about herself; she was able to comply with her work schedule.

“AppInsight allowed me to think about my activities” – 6.0 [1.5] – The results from this statement show that participants were able to identify past computer activities using a small set of attributes, such as application’s name, titles and date. Moreover, it also allowed them to reminisce and think about what they did, which potentiates the adoption of new and improved usage behaviors.

“AppInsight could help me make a better use of my time” – 6.5 [2.0] – For the most part, participants see in *AppInsight* the opportunity to improve their computer usage behaviors. Although the tool was not built to achieve this goal, some participants stated seeing their computer activities could have a positive effect on them. Participant 4 suggested: *“It [AppInsight] could warn me if a specific application is used too much. For instance, I could define that this week I want to play only for two hours. If I played for more than two hours I would like to be notified. That way this [while looking at his usage history] would not happen [smiling]”*.

At the end of the debriefing interview, we asked two last questions directly related to the usefulness of *AppInsight*:

“Would you like this tool to be included in current operating systems?” Nine of our ten participants answered yes to this question. Still, participants showed some privacy concerns and pointed out that users should be aware of the monitoring process and how to disable/uninstall it. Although one participant answered no, he acknowledged that it can be useful for some people. However, in his case, he did not care how much time he spent on his computer or what he did. Nevertheless, this result shows that, indeed, users need this type of tools and wish to visualize their computer usage history. This was once again confirmed when we asked participants *“Would you like to keep this application installed?”* – This was the ultimate usefulness question. Again, nine of ten participants responded yes and kept the application installed on their personal computers.

In the remaining of this section, we present some case studies and participant feedback obtained during the interviews.

Computer usage patterns. While skimming their data, several participants stated that they could clearly see their overall pattern and routine of computer usage. Also, they enjoyed having expectations of their usage confirmed by the visualization while still being able to discover new information. For instance, whilst visualizing his monthly activity a participant commented: *“There is an interesting pattern here [AppInsight]. I often use my computer from 10 a.m. until 2 a.m.; particularly, between 10 a.m. and 7 p.m., and then between 10 p.m. and 2 a.m. This is clearly noticeable here.”*

Additionally, the same participant was able to see how his *work patterns evolved* in the last 3 months, adding: *“It [the work] was much more intense in July. In August it decreased a lot ... oh this is nice, in August my usage is between 3 p.m. and 5 p.m. and between 10 p.m. and 2 a.m. This is when my daughter goes to sleep [laughing]”*

Reminiscing past events. While exploring their data, we notice that participants often made many personal comments as if they were speaking with themselves about past events. Particularly, the list of most used applications often reminded participants about their daily activities and how they evolved over time. One participant said that *AppInsight* was just like a personal diary. Participants enjoyed looking at several types of activities: writing, web browsing, music playing, chatting, and many others. One participant was able to recall what games he played during the holidays and how his gaming pattern evolved: *“I found an interesting pattern. Every time I start to play a new game, I always start big time. I use to spend the first two days playing it most of the time and then it slows down until eventually I stop playing it. I can also see that in the beginning, I always look up for tutorials and FAQs...”*

Lifestyle awareness. Most participants, while looking at *AppInsight*, were able to assess their lifestyle over a specific period of time. For instance, holidays were easily identifiable through both overall usage pattern and type of most used applications: *“Here [pointing], in the end of July I have much less activity ... It was my holidays. The applications also differ; In this period of time I only used my email client and facebook, work applications just disappeared. Two weeks later, it [work] starts again”*. Interestingly, participants were also able to recall non computer based events, such as nights out, birthdays, weddings, dinners with friends and other social events.

Productivity assessment. Although the visualization has shown to be a great way to assess computer usage for all participants, it was especially interesting for those who used their computers for work purposes. Participant 9, who works from home, was very excited to see for the first time her work patterns. Her first comment was: “I’m so happy! Do you see this? I work a lot!” While skimming her data, she added: “I’m so happy that most used applications are work applications. This means that I spend more time working than doing other stuff. I’ve never evaluated it. Since I work from home this is important, because one can lose herself doing other things, right? I was a bit afraid of seeing my computer usage, but now ... I’m so proud of myself.”

The visualization also enabled participants to see periods of intensive work. For instance, one participant was surprised about the amount of time spent at the computer by the end of September. He then remembered that it was due to a paper deadline (see Figure 1): “Wow, there is a lot of usage here ... [drags the timeline’s bar] ... oh yes, this was when I was writing my paper. Intensive work here [pointing to the visualization], why isn’t it always like this?”

Retrospection and improving behaviors. Most participants remarked that *AppInsight* can be a useful tool to evaluate and improve computer usage behaviors: “This allows me to see how much time I’m spending on things I shouldn’t ... It makes you think about it. If I saw that my activity was inappropriate I probably would make an effort to change it”. This was somehow surprising, since *AppInsight* is a visualization tool and does nothing to persuade users to change their behaviors. Nevertheless, some participants were keen to alter their computer usage routine and even gave us some suggestions to improve the visualization: “It would be nice to set an amount of time to play video games, and then it could warn me if that limit was exceeded”.

Detecting unfamiliar activity among personal data. While using *AppInsight*, participants easily related themselves with the visualization. Although they could not remember all their activity, whilst skimming their data they could recall and make sense of it. One participant remarked while trying to explain his computer usage pattern: “... this is so much more obvious to me than to someone else”. Thus, it is no surprise that participants could detect someone else’s usage in their computer. An unknown application or usage pattern was immediately distinguishable. One of the participants that lived with her boyfriend commented: “this is clearly my boyfriend’s usage ... sports, facebook and people I don’t know ... uhm mostly girls ...” Although, her boyfriend was also aware of the monitoring tool, obviously, this comment raised several privacy issues that we need to take into account.

Work tool. During this user study, three of the participants commented that this tool would be useful in their professions. A psychology student stated: “with this I could study people’s behaviors when using their computers. It would be a very useful tool”. Also a marketer suggested: “This can be applied in marketing ... by observing people’s usage patterns and then identify the most appropriate time and content to advertise”. Finally, other participant found it useful to aid her when filling activity reports: “I have to fill reports about my daily activities. I usually use the Outlook to see what mails I replied to, but with this [AppInsight] it is much easier.”

6. DISCUSSION AND LESSONS LEARNED

After analyzing the results we are now able to answer the research questions proposed at the beginning of this user study:

1. *Is AppInsight useful?* Most users found our visualization to be useful. *AppInsight* was especially valuable for those that use computers as a working tool, as it offered a visual and interactive way to monitor their productivity and be aware of possible sources of distraction. This is corroborated by users affirming that our tool allowed them to learn more about themselves and to think more about their activities.

2. *Does AppInsight allow users to reminisce about past events?* We are positive that our technique does indeed offer means to reminisce the past. Users affirmed to be able to recall events from their life that could be related or not to any computer activity, by using *AppInsight*. As an example, we had users recalling days they went to special places, such as the beach; or in which they met or had people over; or even days wherein a relative had specific health problems. Our early results encourage us to explore new contextual cues in order to increase the stimuli conveyed to the users’ memory.

3. *Do participants enjoy using AppInsight?* Nine of our ten users confirmed they would like such a tool to be included on the operating system. Furthermore, every user affirmed to have enjoyed visualizing their computer usage history and they also found the displayed information to be easy to understand. These facts allow us to accept the premise that users actually enjoyed using *AppInsight*.

4. *Would users change their behavior because of AppInsight?* Most users agreed that *AppInsight* would be able to help them make a better use of their time. This highlights the possibility of exploring a more persuasive approach to *AppInsight*’s interface in order to change users’ behavior and promote better working habits and improve productivity.

We learned a few important lessons from this study, which may have some interesting implications for HCI and Visualization:

Less is more. We learned that such a small set of memory cues, which includes the application’s name/URL and window title, allows users to reminisce, and introspect on past events of both their *digital* and *physical* lives. Such capture can be done passively, without disturbing the users, and the resulting data can be used to build more complex information and portray interesting personal trends and patterns.

Computational balance. One of our users complained about computer performance issues during the period in which *AppAware* was installed on his computer. This raises an interesting discussion around the balance between the amount and complexity of the contextual cues and the computational load required to capture them. Although we tried to efficiently explore applicational information, the sheer fact that we adopted a browser independent mechanism for extracting URLs may have caused a slight performance imbalance.

Diversity of uses. Users found several interesting reasons to use *AppInsight*: to help them in their work providing a way to evaluate their productivity and work behaviors; as a memory aid to fill activity reports to show their superiors that they have been working properly; to allow them to think thinking about their own activities and learn something new about themselves; to help them to better manage their time; or just because they enjoy seeing their activity history and how it has evolved over time.

Need for further details. Although users were able to find interesting patterns on their application usage history, our visualization tool did not allow them to explore how time was allocated on each hour, which turned out to be a shortcoming. For example, it was common for users to select an application that

took most of a specific hour and then try to check how the remaining of that hour was spent in the computer, and our interface did not allow that.

Individual behaviors. We believe that users' computer activity has several interesting and barely explored applications. For example, usage patterns could be mined to identify which user is utilizing the computer. This was supported by several users identifying others using their own computers, by just looking into their activity on *AppInsight*. However further research on this matter must be carried out in order to draw any conclusions.

Behaviors may change. Our questionnaires results and participants' comments during the user study demonstrate that they were able to reflect about their activities and were receptive to change their behavior by using our tool. This leads us to believe that *AppInsight* would benefit from some form of persuasive interface, which would help users to change their behavior and eventually improve their productivity.

7. CONCLUSION AND FUTURE WORK

We have presented *AppInsight*, a tool that allows users to visualize their overall patterns of computer usage. By passively monitoring users' activity and collecting two attributes – application's identifier and window title – *AppInsight* was able to effectively characterize users' computer usage behaviors and activities. During our user study, we found that they can derive many different types of benefits from this tool, such as reminiscing past events, being aware of their lifestyle, reflecting about activities, and so on. Overall, users enjoyed using *AppInsight*, so it was that they kept the tool installed afterwards.

Furthermore, we found that *AppInsight* can potentially be used to motivate users to have "better" computer usage behaviors. As future work we intend to persuade users to adopt more appropriate behaviors according to their personal goals and expectations. While we acknowledge that users may benefit from this type of tool, we also know that they usually do not like to be bothered. Therefore, developing such a tool, without annoying users, is a true challenge. Finally, exploring different contextual cues is also an interesting topic for future research.

8. ACKNOWLEDGMENTS

This work was supported by FCT (INESC-ID multiannual funding) through PIDDAC Program funds and the project PAELife, reference AAL/0014/2009. Barata and Nicolau were supported by FCT, grants SFRH/BD/72735/2010 and SFRH/BD/46748/2008.

9. REFERENCES

[1] M. 'Aquino, S. Elahi, and E. Motta. Personal monitoring of web information exchange: Towards web lifelogging. In *Proc. of WebSci10*, 2010.

[2] D. Baur, F. Seiffert, M. Sedlmair, and S. Boring. The streams of our lives: Visualizing listening histories in context. *Visualization and Computer Graphics, IEEE Transactions on*, 16(6):1119–1128, nov.-dec. 2010.

[3] T. Beauvisage. Computer usage in daily life. In *Proc. of CHI '09*, pages 575–584, 2009.

[4] R. A. Bjork and A. Richardson-Klavehn. *On the puzzling relationship between environmental context and human memory*. Lawrence Erlbaum Associates, Publishers, 1989.

[5] R. Boardman and M. A. Sasse. "Stuff goes into the computer and doesn't come out": a cross-tool study of personal information management. In *Proc. of CHI '04*, pages 583–590, 2004.

[6] E. Chi. Improving web usability through visualization. *Internet Computing, IEEE*, 6(2):64–71, mar/apr 2002.

[7] S. Dumais, E. Cutrell, J. Cadiz, G. Jancke, R. Sarin, and D. C. Robbins. Stuff i've seen: a system for personal information retrieval and re-use. In *Proc. of SIGIR '03*, pages 72–79, 2003.

[8] S. G. Eick. Visualizing online activity. *Commun. ACM*, 44:45–50, August 2001.

[9] J. Gemmell, G. Bell, and R. Lueder. Mylifebits: a personal database for everything. *Commun. ACM*, 49:88–95, January 2006.

[10] A. Graham, H. Garcia-Molina, A. Paepcke, and T. Winograd. Time as essence for photo browsing through personal digital libraries. In *Proc. of JCDL '02*, pages 326–335, 2002.

[11] J. Hailpern, N. Jitkoff, A. Warr, K. Karahalios, R. Seseck, and N. Shkrob. Youpivot: improving recall with contextual search. In *Proc. of CHI '11*, pages 1521–1530, 2011.

[12] S. Hangal, S. L. Lam, and J. Heer. Muse: Reviving memories using email archives. In *Proc. of UIST '11*, 2011.

[13] D. F. Huynh, S. M. Drucker, P. Baudisch, and C. Wong. Time quilt: scaling up zoomable photo browsers for large, unstructured photo collections. In *Proc. of CHI EA '05*, pages 1937–1940, 2005.

[14] O. Laadan, R. A. Baratto, D. B. Phung, S. Potter, and J. Nieh. Dejerview: a personal virtual computer recorder. In *Proc. of SOSP '07*, pages 279–292, 2007.

[15] S. Tak and A. Cockburn. Window watcher: a visualisation tool for understanding windowing activities. In *Proc. of OZCHI '09*, pages 241–248, 2009.

[16] M. Torrens and J. Iluís Arcos. Visualizing and exploring personal music libraries. In *Proc. of ISMIR '04*, pages 421–424, 2004.

[17] E. Tulving, W. Donaldson, G. Bower, and U. S. O. of Naval Research. *Organization of memory*. Academic Press, 1972.

[18] F. Viegas, D. Boyd, D. Nguyen, J. Potter, and J. Donath. Digital artifacts for remembering and storytelling: posthistory and social network fragments. In *Proc. of the Hawaii International Conference on System Sciences*, page 10 pp., jan. 2004.

[19] F. Viegas and M. Smith. Newsgroup crowds and authorlines: visualizing the activity of individuals in conversational cyberspaces. In *Proc. of the Hawaii International Conference on System Sciences*, page 10 pp., jan. 2004.

[20] F. B. Viégas, S. Golder, and J. Donath. Visualizing email content: portraying relationships from conversational histories. In *Proc. of CHI '06*, pages 979–988, 2006.

[21] F. B. Viégas, M. Wattenberg, and K. Dave. Studying cooperation and conflict between authors with history flow visualizations. In *Proc. of CHI '04*, pages 575–582, 2004.

[22] S. J. Waterson, J. I. Hong, T. Sohn, J. A. Landay, J. Heer, and T. Matthews. What did they do? understanding clickstreams with the webquilt visualization system. In *Proc of AVI '02*, pages 94–102, 2002.

[23] R. Xiong and J. Donath. Peoplegarden: creating data portraits for users. In *Proc. of UIST '99*, pages 37–44, 1999