# Event Segmentation in the Wild

Brendan Jonesrebandt
Department of Cognitive Science
Distributed Cognition and Human Computer Interaction
University of California, San Diego
La Jolla, CA 92092
Advisor: Professor Jim Hollan

## Introduction

People's everyday lives are filled with a variety of activities. Many times these activities can be broken up into distinct events such as cooking a meal or driving to work. It has been suggested that people spontaneously segment activities into meaningful parts and subparts, which has been supported by behavioral and neuroimaging data [4]. When it comes to remembering the details of these events, new data indicates that those individuals who are better able to segment the activity into events are better able to remember it [15]. Newtson devised a procedure which has been the foundation for many studies in event segmentation. Participants are shown a video of an activity which they are asked to break up into different portions. The participants push a button to whenever they judge that one meaningful event has ended and another one is beginning. They only push the button once to indicate the separation between the two portions of the activity. The activities are usually highly stylized and singly focused activities such as an actor making a bed, washing a car, or setting up a tent. Event segments produced by this task are reliable across viewers [16]. These event boundaries, the points judged to be between two portions of the activity, typically are described as correspond to sub-goals of an actor, such as the stages in the Figure 1[5]. In the example, the main goal of the activity is setting up the tent, but users found six sub-goals within the overarching main goal. The sub-goals might be categorized as the following by participants; a) unpack the tent, b) lay out the tent, c) insert the poles, d) erect the tent, e) put on the rain fly, f) finished.

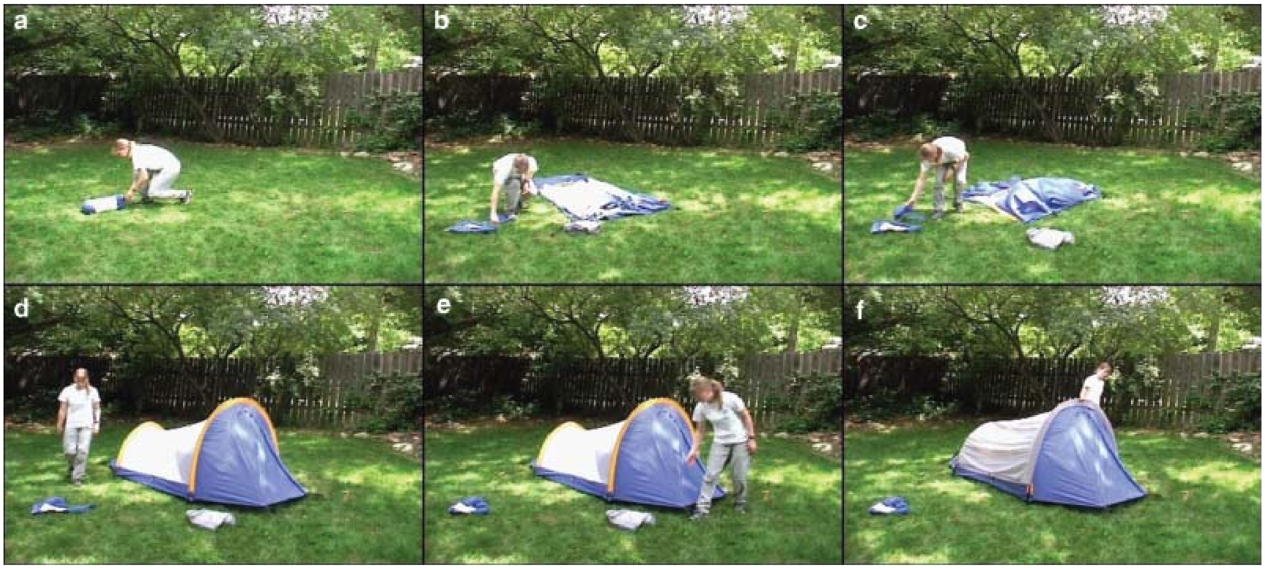

**Figure 1**: Example of event boundaries. Illustrating different phases of setting up a tent and how users split the activity into segments [5].

Event segmentation has been suggested as a method to organize the large, complex data sets being created by life-logging technologies [8]. Life-logging, a term coined by Gordan Bell of Microsoft Research, refers to the practice of digitizing all of the information produced by an individual in everyday life. Life-logging methods provide an avenue for individuals to track information that is important to them. Studies have focused on designs and prototypes of tools to assist with recall [17, 23]. Currently, many people engage in life-logging in the form of creating and maintaining blogs [1, 2]. Some emerging technologies focus on a newer form of life-logging, visual life-logging, which is based on passive image capture of a person's experiences [1]. These emerging technologies can be used for a wide variety of purposes such as monitoring sleep and exercise patterns, recording vacations, and assisting individuals with short term memory loss.

One of the main problems with visual life-logging is the huge collection of complex data that arises. Often the images that are captured are paired with other forms of data such as audio, temperature, light levels, and movement. In order to use this data, an effective way to view it must be implemented. A method currently used to organize these large sets of complex data is event segmentation. This method of organization, which may be similar to the way people naturally organize activity, is the process by which individuals divide a continuous stream of activity into meaningful segments [5]. In addition, by breaking up the data into meaningful chunks, the data is grouped by activity rather than by time.

This approach for organizing the data has a few possible problems. One problem is that many natural activities appear to be very broken up, which may create too many segments to be useful. Everyday activities are full of random encounters with other individuals, miscellaneous phone calls, emails, and time constraining events, as well as a plethora of other actions that may not be part of a certain event, but are just a part of life. These "interruptions" in an activity may make it more difficult to meaningful divide an event up into a series of segments, as some studies indicate. One study found that on average, mobile professionals received just over four interruptions per hour [23]. This study found that in 64% of the interruptions, individual received a benefit from the interruption, though 40% of the time the individual did not resume the work they were doing prior to the interruption [23]. A field study conducted observations of workers in the roles of analysts, software developers, and managers [19]. The study found that there was a high level of discontinuity in the execution of activities, with an average of three minutes on a task before switching to a different task [19]. These studies analyzed everyday activity as opposed to the previous work done with highly stylized activities. The intent of the segmentor can also change how an activity is segmented. In another experiment, participants were asked to segment a video using Newtson's method while performing one of the following tasks: learning the task being performed by the actor or forming an impression of the actor's personality [20]. It was found that there was poor within-participant agreement on the location of unit boundaries between the two conditions [20]. Participants who were focused on the impression of the actor segmented different sections than those that were focused on learning the task.

# Rational of Research Project

If life-logging is to be useful, there must be an efficient and effective means of storing and sorting through the large, complex data sets which arise from the life-logging. With massive amounts of data, being able to find specific things becomes very important. If segmentation is the method of organizing these data sets, a topic that has received little study, but is of paramount importance is how individuals segment their own data. This preliminary study analyzes the difference between segmenting real activity as opposed to highly stylized events. Findings from event segmentation in a controlled lab may have poor external validity when compared to segmentation of activities that occur in daily life. In addition, the difference between a participant segmenting their own data compared to someone else segmenting their data will be investigated. One thought is that the more complex the activity, the less likely participants, other than the one who recorded the data, will be able to make sense of it. Segmenting an individual setting up a tent is a fairly simple, straightforward task. Segmenting a task that is unfamiliar to a participant and full of many people and actions is much more complex and most likely much more difficult to segment. The other participants will likely see this complex activity as one filled with many interruptions just as the observers in the previously mentioned study saw many interruptions during tasks [19, 23]. A key question though is whether the individual will see their activity as full of interruptions. It will also become clear through our data whether participants are able to remember and utilize information for segmenting their activities that is not apparent in the data. These non-visible contextual clues will be due to context reinstatement. Context reinstatement is the reestablishment of relevant facts and circumstances surrounding an event, many times bringing back memories and the thoughts that occurred in that environment. Evidence has indicated that the mind and brain are tracking features of an individual's environment and create event boundaries when a salient feature changes unpredictably [5]. If this method of creating event boundaries is employed by participants, there may be cases when individuals will create segments based off of these salient changes which are not visible to the participants who did not record the data. This may further increase the disparity between segments between the participants who are segmenting another individual's data compared to the individual segmenting their data. This study will analyze the reliability of segmentation between individuals on real, everyday activities and any changes that occur from a participant segmenting their own data compared to others segmenting the data.

# Life-Logging Tools

In order to get a clearer picture of how individuals segment their own data, we decided it was essential to sample different types of data to determine whether there were differences between the different types of data in addition to differences between participants segmenting their own data sets versus someone else's. We decided to focus on three different types of life-logging data: pictures, hand written data, and screenshots of computers. This would allow us to not only analyze different types of data, but different types of activities as well. Pictures were collected utilizing the Microsoft SenseCam, hand written data was captured with the Livescribe Pulse Smart Pen, and computer screenshots were gathered with the ActivityTrailsLite program. These particular devices were chosen because they are non-invasive and easy to use and operate. The SenseCam is simply worn around the neck and automatically takes pictures. The Livescribe pen is similarly easy to use and was able to easily replace a traditional pen or pencil for a written based activity. ActivityTrailsLite runs in the background and performs its screenshot capturing without disrupting the user.

## Microsoft SenseCam

The SenseCam (see **Figure 2**) is a wearable, automatic digital camera developed by Microsoft research [9].

**Figure 2**: The Microsoft SenseCam

This SenseCam was originally developed as an aid for people with memory loss [9]. It has been shown to help people reconnect with past activities [22]. In particular, it was shown to help patients with amnesia recall events [21]. The SenseCam has expanded to a tool tested in a wide variety of fields including tourism, patient care, education, and accessibility within business [3]. The SenseCam contains a variety of sensors including a light-intensity sensor, light color sensor, passive infrared sensor, temperature sensor, and a tri axis accelerometer [9]. When any of the sensors are triggered, such as by a change in temperature, light, or movement, the device is configured to automatically take a photograph [9]. This camera is intended to be worn around the neck and is constructed with a wide-angle lens in order to better capture the scene and view of the wearer [9]. The quality of the pictures varies, though they are fairly discernable for the most part if the scene is well lit (see **Figure 3**). Though it automatically takes pictures, there is also a button for manually taking pictures in addition to a button which stops the camera for a short time period. The pictures are taken at VGA resolution (640x480 pixels). The SenseCam can store approximately 30,000 of these compressed .jpg images [9]. The data from the sensors, as well as these pictures, are time stamped when stored [9], allowing their parallel streams of data to be easily matched. At this time no audio detection or recording is included although the device does emit certain scripted sounds in order to provide some feedback to the user [9].


**Figure 3**: Example of a SenseCam image quality.

## Livescribe Pulse Smart Pen

Another tool we used in order to collect visual life-logging data was the Livescribe Pulse smart pen (see **Figure 4**).


**Figure 4**: The pen interacting with the digital paper interface. It is currently poised on the record action which will allow it to record the audio information in sync with written notes.

Pens and "digital" paper have become much more common and provide a medium to transfer written information to a digital format. There have been several forays into using these tools in a variety of situations such as 3D model design [26], interactive pop-up books [27], and linking digital and paper copies for manipulation through a cell phone [28]. Because of these advances in technology, it is possible to capture written information in a digital format. The Livescribe pen can be used to record synchronized audio and written information when it is used with special digital paper [10]. The digital paper uses a series of dots on its surface that are used to define the location of the text. The paper then acts as an interface to as access different sections of the audio notes by taping the appropriate area of the text. This tool allowed us to collect written data, which was composed primarily of pictures and words.


**Figure 5**: Example of data collected by Livescribe pen

# ActivityTrailsLite

The third visual life-logging technology we used was ActivityTrailsLite. This software is a basic MAC OSX application that takes screenshots of the user's computer screen every five seconds (see **Figure 6**). At this time there is no other corresponding data taken. The data captured is stored as a series of images which can be viewed with a variety of tools. This presented a third medium to analyze the similarity of event boundaries between participants.

Icon indicating ActivityTrailsLite is running

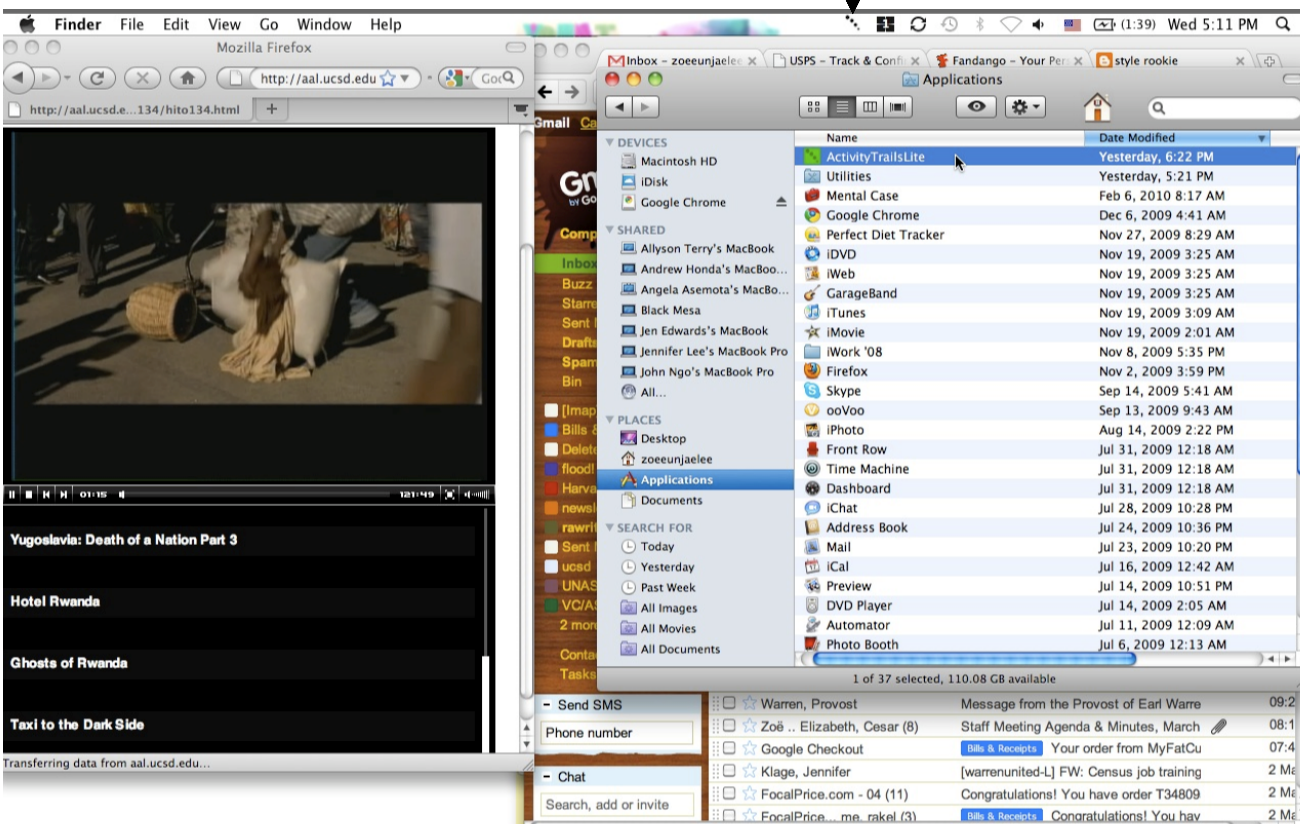


**Figure 6**: Sample screenshot taken by ActivityTrialsLite. The icon on the top toolbar is the only indication that the software is even running, therefore it does not alter the way a user performs tasks in the slightest.

# SenseCamPresenter

The SenseCamPresenter, developed by the DCog-HCI Lab at UCSD, has been the main tool used for participant event segmentation. The software was developed for displaying the images from all three life-logging tools in an easy to view interface (see **Figure 7**). A data set collected through the use of a life-logging tool is loaded into the SenseCamPresenter. The software then displays the visual components of the captured data (pictures from the SenseCam, screenshots from ActivityTrailsLite, or images of the text and/or pictures from the Livescribe paper and pen). Users then navigated through this data and segmenting it.

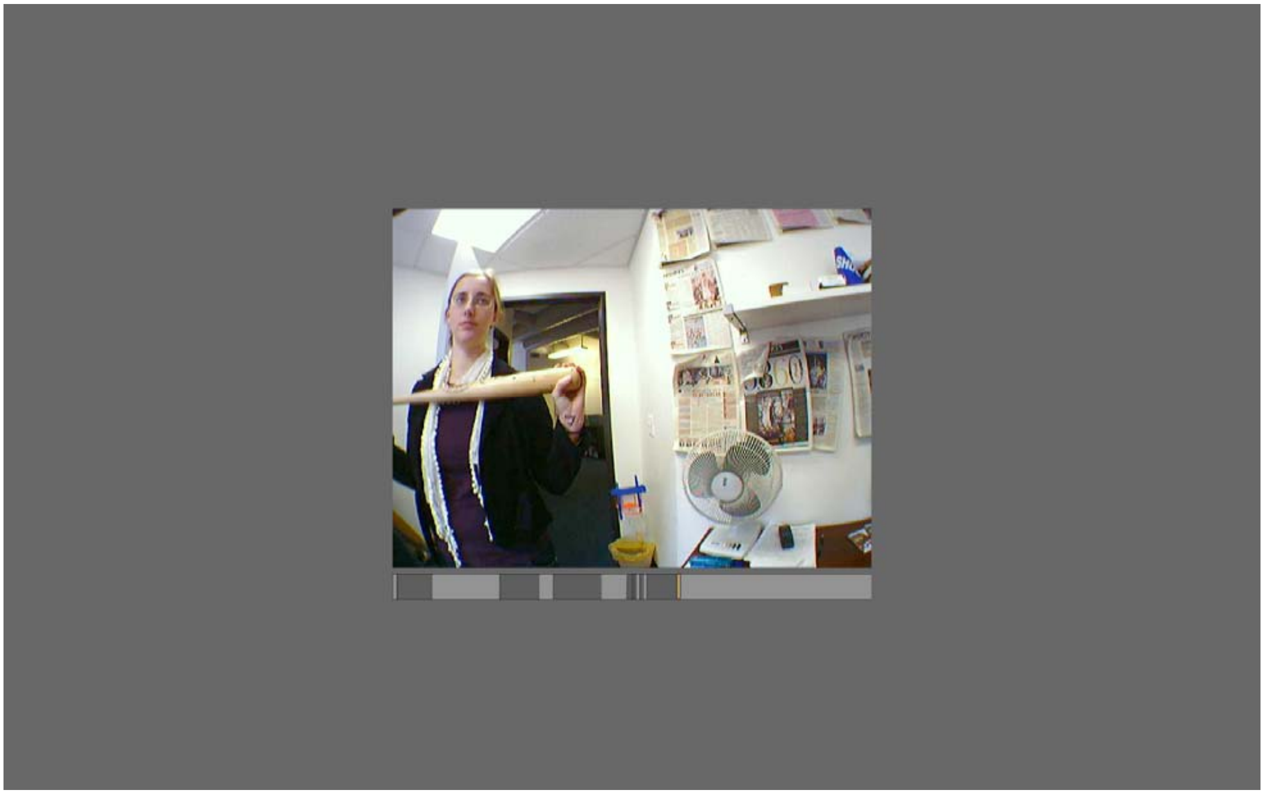

**Figure 7:** A screenshot of the SenseCamPresenter displaying a section of SenseCam data partially segmented. The yellow bar indicates where the user currently is in the dataset (the picture currently being viewed). The different shades of gray indicate different event segments that the user has divided the data into.

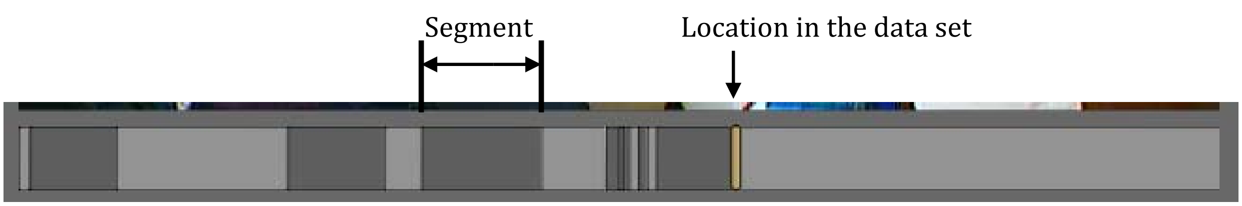


**Figure 8**: A close up of the bar indicating segments and the location in the data set.

The SenseCamPresenter works in conjunction with a dial in order to navigate through the data and mark event boundaries (see **Figure 9**). The dial has five programmable buttons which were all set with the same function, creating an event boundary. The participant could use any or all of the buttons, since they had they were programmed for the same function. The inner ring of the dial, or “jog”, rotates through 360 degrees and displays the pictures frame by frame. The outer black ring, or “shuttle” was rubberized and spring loaded. It was used to fast forward or rewind; the more you turn it clockwise or counterclockwise, the faster it fast forwards or rewinds, respectively.

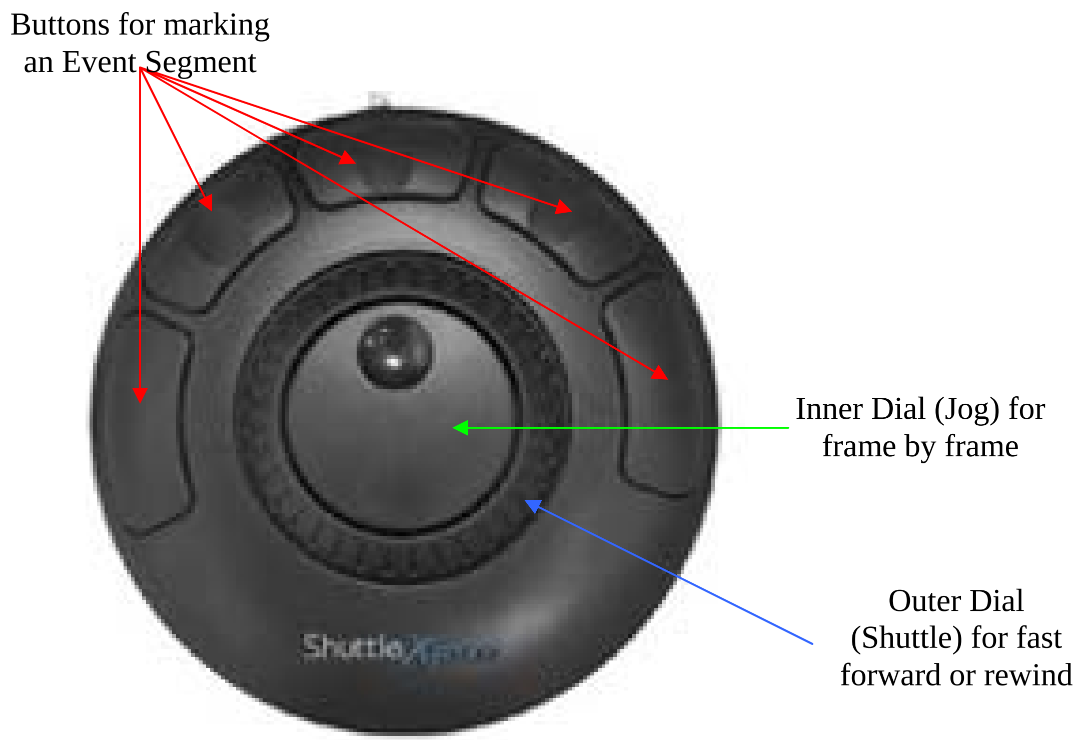


**Figure 9**: The dial interface for controlling the SenseCamPresenter

# Methods

## Overview

This preliminary study analyzed participant's event segmentation of real life activities, paying special attention to the differences between participants's segmenting their own data compared to this same data being segmented by others. Participants recorded an activity with one or more of the three recording devices: the Microsoft SenseCam, ActivityTrailsLite, or the Livescribe Pen. They were then asked to segment their data set, as well as up to three other participant's data sets, utilizing the SenseCamPresenter interface. The data sets varied in length and content. The interface used to segment the data sets also recorded the participant's interaction with the data and where their segments occurred. There were four weeks of pilot studies to fix any bugs and fine tune the procedure.

## Participants

Participants were 17 undergraduate students at the University of California, San Diego (9 Males and 8 Females). Due to the nature of this study, finding participants who were trustworthy and able to keep the recording devices safe while in their possession was of paramount importance. A list of 30 possible participants was formulated from students and acquaintances at the University of California, San Diego. These participants spanned a range of 18-23 years of age. The participants were interviewed to determine whether they were acceptable for our study. The questions asked during the interviews were designed to help us choose an activity that the individual already engaged in that would provide a good data set. We were looking for activities that spanned one to three hours in order to ensure that there was enough time to provide several different segmentation points, but not too long that it would discourage participants from carefully reviewing the data set. In addition, participants were also asked about their comfort level with recording their activities and being available to review their data set. This survey helped us eliminated any participants who were not comfortable with recording their activities, unable to participant due to time constraints, or lacked appropriate activities.

## Recording Devices

The three recording devices consisted of the Microsoft SenseCam, ActivityTrailsLite, and the Livescribe Pen. The device chosen for the participant depended on the activity that was to be recorded. For computer based activities, ActivityTrailsLite was utilized. For writing or drawing activities, the Livescribe Pen was chosen. For physical activities, the Microsoft SenseCam was chosen. There were 5 participants who recorded activities recorded with ActivityTrailsLite (2 Males and 3 Females), 2 participants who recorded with the Livescribe Pen (1 Male and 1 Female), 8 Participants who recorded with the SenseCam (5 Males and 3 Females), and one participant who recorded with both ActivityTrailsLite and the SenseCam simultaneously.

## Activities Recorded

In order to test a wide range of activities, we utilized the survey questions to find activities that engaged in that were different from any that we had already recorded. ActivityTrailsLite recorded a participant filling out a job application, writing a newsletter, playing a video game, watching a movie while taking notes on it, and writing an essay. The SenseCam recorded a participant moving between and in a class at the university, working at an after school education center, eating dinner, a hiring workshop, a large cultural program, a party, a choir performance, a study session, and an editor working for a newspaper. The Livescribe Pen recorded a participant drawing and another one taking notes in class. One of our participants used both the SenseCam and ActivityTrailsLite to record working on homework from two different perspectives.

## Procedure

Once it was determined what activity each participant was going to record, the appropriate life-logging tool was selected for them to use. Each participant was trained only in the use of the recording device that was going to be used for their activity. Training occurred up to three days before the activity to be recorded took place. Participants who were using ActivityTrialsLite were shown how to download the software from the email they received with the application attached. They were shown how to start and stop the application as well as how to check to make sure that it was working. Participants were also shown how to transfer the images through the use of a USB drive. Participants who were using the Livescribe pen were instructed how to turn it on, how to use the dot-patterned paper to record what they were using the pen for, and how to turn off the pen. We had them demonstrate how they were going to use the pen to ensure that they understood the instructions and would collect data in an appropriate fashion. Participants who were using the SenseCam were instructed how to wear it, how to turn it on, how to know when it was operating, how to manually take a picture, how to use the do-not disturb button (pauses the SenseCam for five minutes), and how to turn it off. We had them test all of the functions and practice turning it on and off as well as recognizing when it was operating. Because the SenseCam was worn in public and could possibly cause questions and/or concerns, users were also given a card with instructions of how to respond if anyone approached them:

> *"I am participating in an experiment on everyday memory. This is a digital camera that automatically captures low-resolution still images throughout the day, which will later be used to test my memory. It does not record audio or full-motion video. Any images captured will not be made public in any fashion and will only be seen by myself, during the memory tests, and by the experimenters. If you would prefer, I can turn off or temporarily deactivate the camera, and/or make a note and have the images just taken deleted without anyone seeing them. I can also provide the contact information of the experimenters."*

In addition, each participant was given a user guide that repeated all of the instructions and had our contact information available in case they ran into difficulties or had questions.

The participant then recorded their selected activity which ranged from 1 to 3 hours. After the activity was recorded, the data was transferred to a computer in our lab by USB for ActivityTrailsLite or directly from the device for the SenseCam and Livescribe pen. This transfer was done within two days of recording the activity so that it could be reviewed with the SenseCamPresenter to ensure the device recorded correctly and the data was usable (see **Figure 10**)

**Figure 10**: Testing Livescribe data with the SenseCamPresenter

Participant's reviewed their data approximately one week after recording it. The participants interacted with the sets of data in a small room in front of a computer terminal with the SenseCamPresenter (see **Figure 11**). Each session was done with a single participant. The participants were told that they would be audio and video recorded in order to create a record of their interaction with the data that could be reviewed at a later date. The experimenters also passively observed and were available if there were questions or difficulties operating the interface. The participant's interacted with their data through the dial interface described earlier.

**Figure 11**: Participant segmenting their data set

Participants were instructed how to operate the dial interface in order to review their data and were told to “talk about anything that comes to mind while viewing the data.” This general prompt left very little constraints on the participant, allowing them to talk about anything that they remembered. Having an individual verbalize about a recording of their own activity is known as Individual Auto-confrontation [29]. Participants auto-confronted their recordings by using either of the two dials, allowing them to move through the data at the speed they preferred. The buttons were disabled during this portion of the experiment. Once the participants completed auto-confronting their data, they were asked to perform the second part of the procedure, segmenting the data set. Participants were instructed to “break the data up into meaningful pieces.” This general set of instructions ensured that no participant’s were biased in any way in how they decided to segment the event. Participants were shown how to use the dial interface to place an event boundary, which marked the end of one segment and the beginning of the next. The dials on the device performed the same as in the previous portion of the experiment. As the participants went through their data, they were asked to describe why they placed an event boundary each time they marked one. This was audio and video recorded.

After the participants completed segmenting their own data set, they were asked to perform the same segmentation activity on 1-3 other data sets depending on how many were available at the time. They were free to look through the data first or just began segmenting. They were given the same instructions: "break the data up into meaningful pieces." The goal was to have at least 3 participants segment each data set in addition to the owner of the data set. This provided three sets of event boundaries to compare to the original participant's event boundaries in addition to each other's. After completing the last data set, participants were told what we were studying and we answered any questions that they had as well as took feedback on any difficulties or improvements that could be made.

Once the participants had completed all of their segmentations, the video log was analyzed and their reasons for segmenting and the number of segments they made for each data set was coded. Once all of the participants' information was coded, the reason for their segment was broken down into five categories: Scenario Cue, Visual Cue, Natural Break, Contextual Cue, or Other. The data was sorted into the appropriate category based on what was captured by the video and audio recordings. Scenario cues were instances where the participant created a hypothetical situation or relevant experience they had to understand what the individual was doing. Visual cues were cases where the participant focused on changes in the environment, such as location, to determine when to segment. Natural Breaks were culturally seen as stopping points in an activity, such as going to the bathroom or going on Facebook after working on a paper. Contextual Cues were segments where the reason for the segment could not be seen, such as the participant remembered a conversation well they were on a certain layout on their computer screen. The other category contains the segments that do not fit into these categories. Sometimes participants would segment because they felt they should, but they could not give a better reason than that; other times they did not say why they segmented. After all of the segments were coded for their type of segment, they were organized into tables (see **Figures 12-15**).

# Results

Of the original 17 participants, there was at least partial data corruption on 3 participants, rendering either their video data (their reasoning behind their segments) inaccessible, their actual segmentation data (the exact placement of segments) inaccessible, or both. Of the other 14 participants we found that their reasoning for event boundaries fell into 5 categories: Scenario Cue, Visual Cue, Natural Break, Contextual Cue, or Other. Participants utilized one of two methods to begin segmenting: initially looking through the data and then going back and segmenting, or just starting out segmenting from the get go. Both methods were used about equally. It seemed to be based more on an individual preference rather than the data, since the people that used this strategy would usually use it on every data set they segmented and those that did not would never use it.

| **Participant** | **Scenario** | **Visual** | **Natural Break** | **Contextual** | **Other** | **Total** |
|---|---|---|---|---|---|---|
| S | 13 | 0 | 2 | 3 | 0 | 18 |
| E | 7 | 0 | 10 | 0 | 0 | 17 |
| U | 15 | 42 | 5 | 9 | 0 | 71 |
| K | 12 | 2 | 9 | 1 | 2 | 24 |
| I | 13 | 10 | 1 | 1 | 2 | 25 |
| L | 71 | 28 | 7 | 5 | 0 | 111 |
| B | 6 | 3 | 2 | 0 | 3 | 11 |
| T | 18 | 6 | 2 | 1 | 7 | 27 |
| Z | 33 | 5 | 4 | 13 | 0 | 55 |
| D | 28 | 9 | 3 | 2 | 4 | 42 |
| P | 33 | 26 | 3 | 9 | 3 | 71 |
| J | 63 | 57 | 4 | 7 | 2 | 131 |
| H | 23 | 8 | 5 | 0 | 3 | 36 |
| O | 15 | 13 | 14 | 0 | 2 | 42 |
| **Total** | 350 | 209 | 71 | 51 | 28 | 681 |
| **Percentage** | 0.513950073 | 0.306901615 | 0.104258443 | 0.074889868 | 0.041116006 | 1 |

**Figure 12**: Event Segment by participant by type of segment. The total number of segments was 681 between the 14 participants. 51% were Scenario, 31% were Visual, 10% were Natural Break, 7.5% were Contextual, and 4% were classified as other.

The Scenario Cue category encompassed approximately 51% of the event boundaries (see **Figure 12**). This category was defined by participants creating or describing a scenario that attempted to make sense of the current actions and activities, allowing them to break them into meaningful chunks that fit within these scenarios. Some scenarios that users created were going to a lecture or creating a newsletter. Many times the scenarios were proven to be inaccurate later in the data set, but the participant usually kept their scenario and their event boundaries. Because the participants segmented their own data sets and knew what the exact scenario was, this category is most likely a larger percentage than it would be if no one segmented their own data sets. The variability within this category between data sets was quite large. The lowest and highest percentages in a full participant comparison (4 participants segmented the same data set) was 10% and 71% respectively for the entire data set.

The Visual Cue category made up about 31% of the event boundaries (see **Figure 12**). This category of event segmentation was employed by participants when they would focus on a visual cue, often location. This occurred primarily when participants were unable to create a scenario to fit the activities that were occurring, typical of the more complex, irregular activities such as working at an afterschool help center and playing an online video game. Both are activities that do not necessarily have a rigid structure and order to them if one is unfamiliar with them. When using the Visual Cue, participants would often lock onto one type of change, such as going in or out of buildings. This often caused them to create a large number of event boundaries in comparison to the participant who recorded the activity and used scenarios to create event boundaries. In addition, very few participants had a significant number of event segments falling under this category when they were segmenting their own data. This category ranged from 0% to 77% depending on the data set.

The Natural Break category was about 10% of the event boundaries (see **Figure 12**). This category was similar to the Visual Cue category, but was more culturally recognized as a natural break point. These included things such as saving a word document, taking a bathroom break, or ending a page of notes. The Natural Break category also differed from the Visual Cue category in how participants identified it as a segment; there were much fewer Natural Breaks than either Visual Cues or Scenario Cues. Those that were utilizing scenarios or visual cues to help them segment had the most agreement with their segments when there was a Natural Break. Natural breaks were more common in less complex activities. 0% to 65% of the event segments of a data set were natural breaks.

The Contextual Cue category consisted of 7.5% of the event boundaries (see **Figure 12**). Contextual Cues were things not visible in the data set which the segmentor based their segment off of. Interestingly enough, there were a few instances where one of the participants recognized the activity or event and was able to create a couple of segments that fit into this category, even though they did not record the event themselves. Even thought this category is a small percentage overall, several participants had many segments that fell under this category, one even having 50% of his segments for his data set falling in this category. There were between 0% and 18% contextual cue event boundaries per data set.

The Other category consisted of 4% of the event boundaries (see **Figure 12**). The other category contains the segments that do not fit into the other four categories. Sometimes participants would segment because they felt they should, but they could not give a better reason than that; other times they did not say why they segmented-both of those conditions fell under this category. The Other event boundaries per data set comprised 0% to 29%.

In addition to analyzing based on an individual bases, participants were compared against each other for the same data set as well as comparing the same types of data sets (SenseCam compared to SenseCam and so forth), to analyze how accurate segments are between individuals and devices. As is visible in Figures 13-15, there are no clear indications of the type of segment preferred for each type of device. The variability between individuals and between and within devices is very large. In addition, the number of segments an individual makes per data set varies drastically, even within data sets. The average agreement between at least two individuals is only 13%. The average agreement between all four participants is slightly less than 3%.

The average number of segments per data set is a little bit less than 50 and the range for a data set goes from 20 to 134. Within a data set the largest range is 8 to 55 segments. There is a huge amount of variability between individuals and between and within devices. To get a clearer picture of the difference between individuals when segmenting the same data set see **Figures 16-21.**

| | | | **Livescribe** | | | |
|---|---|---|---|---|---|---|
| **Participants** | **Scenario** | **Visual** | **Natural** | **Other** | **Contextual** | **Totals** |
| **P** | 9 | 0 | 0 | 0 | 9 | 18 |
| U | 1 | 0 | 4 | 0 | 0 | 5 |
| J | 7 | 14 | 0 | 1 | 0 | 22 |
| S | 0 | 0 | 5 | 0 | 0 | 5 |
| **Total** | 17 | 14 | 9 | 1 | 9 | 50 |
| **Percentage** | 0.34 | 0.28 | 0.18 | 0.02 | 0.18 | |
| | | | | | | 0 |
| **U** | 0 | 11 | 0 | 0 | 0 | 11 |
| L | 4 | 27 | 2 | 0 | 0 | 33 |
| J | 2 | 7 | 4 | 0 | 0 | 13 |
| O | 0 | 1 | 2 | 0 | 0 | 3 |
| **Total** | 6 | 46 | 8 | 0 | 0 | 60 |
| **Percentage** | 0.1 | 0.77 | 0.13 | 0 | 0 | |

**Figure 13**: Comparison of the Livescribe data sets. The first participant listed is the one that recorded the data set.

| | | | **SenseCam** | | | |
|---|---|---|---|---|---|---|
| **Participants** | **S** | **V** | **N** | **O** | **C** | **Totals** |
| **A** | 0 | 0 | 5 | 2 | 1 | 8 |
| K | 0 | 0 | 4 | 0 | 0 | 4 |
| B | 2 | 0 | 2 | 0 | 0 | 4 |
| T | 2 | 0 | 2 | 0 | 0 | 4 |
| **Total** | 4 | 0 | 13 | 2 | 1 | 20 |
| **Percentage** | 0.2 | 0 | 0.65 | 0.1 | 0.05 | |
| | | | | | | |
| **K** | 6 | 2 | 0 | 1 | 1 | 10 |
| T | 11 | 6 | 0 | 7 | 0 | 24 |
| A | 8 | 6 | 0 | 3 | 0 | 17 |
| **Total** | 25 | 14 | 0 | 11 | 1 | 51 |

| | | | | | | |
|---|---|---|---|---|---|---|
| **Percentage** | 0.49 | 0.27 | 0 | 0.22 | 0.02 | |
| | | | | | | |
| **B** | 3 | 0 | 2 | 1 | 0 | 6 |
| Z | 4 | 3 | 2 | 0 | 0 | 9 |
| D | 7 | 1 | 1 | 0 | 0 | 9 |
| U | 6 | 18 | 0 | 9 | 0 | 33 |
| **Total** | 20 | 22 | 5 | 10 | 0 | 57 |
| **Percentage** | 0.35 | 0.39 | 0.09 | 0.18 | 0 | |
| | | | | | | |
| **T** | 5 | 0 | 0 | 0 | 1 | 6 |
| Z | 5 | 2 | 1 | 0 | 1 | 9 |
| D | 9 | 2 | 1 | 0 | 0 | 12 |
| P | 8 | 3 | 3 | 0 | 0 | 14 |
| **Total** | 27 | 7 | 5 | 0 | 2 | 41 |
| **Percentage** | 0.66 | 0.17 | 0.12 | 0 | 0.05 | |
| | | | | | | |
| **Z** | 22 | 0 | 1 | 0 | 11 | 34 |
| P | 10 | 21 | 0 | 2 | 0 | 33 |
| J | 18 | 26 | 1 | 0 | 0 | 45 |
| S | 12 | 6 | 4 | 0 | 0 | 22 |
| **Total** | 62 | 53 | 6 | 2 | 11 | 134 |
| **Percentage** | 0.46 | 0.4 | 0.04 | 0.01 | 0.08 | |
| | | | | | | |
| **D** | 4 | 0 | 1 | 1 | 2 | 8 |
| P | 8 | 3 | 3 | 0 | 0 | 14 |
| U | 8 | 12 | 1 | 0 | 0 | 21 |
| S | 7 | 2 | 0 | 1 | 0 | 10 |
| **Total** | 27 | 17 | 5 | 2 | 2 | 53 |
| **Percentage** | 0.51 | 0.32 | 0.09 | 0.04 | 0.04 | |
| | | | | | | |
| **L** | 9 | 0 | 2 | 0 | 5 | 16 |
| M | 14 | 0 | 0 | 0 | 0 | 14 |
| **Total** | 23.5 | 0.32 | 2.09 | 0.04 | 5.04 | 30 |
| **Percentage** | 0.78 | 0.01 | 0.07 | 0 | 0.17 | |
| | | | | | | |
| **H** | 4 | 0 | 0 | 0 | 0 | 4 |
| O | 0 | 5 | 1 | 2 | 0 | 8 |
| I | 3 | 2 | 0 | 0 | 0 | 5 |
| **Total** | 7 | 7 | 1 | 2 | 0 | 17 |
| **Percentage** | 0.41 | 0.41 | 0.06 | 0.12 | 0 | |

**Figure 14**: Comparison of the SenseCam data sets. The first participant listed is the one that recorded the data set.

| | | | **ActivityTrailsLite** | | | |
|---|---|---|---|---|---|---|
| **Participants** | **S** | **V** | **N** | **O** | **C** | **Totals** |
| **S** | 13 | 0 | 2 | | 3 | 18 |

| | | | | | | |
|---|---|---|---|---|---|---|
| E | 1 | 0 | 2 | 0 | 0 | 3 |
| K | 5 | 1 | 0 | 1 | 1 | 8 |
| **Total** | 19 | 1 | 4 | 1 | 4 | 29 |
| **Percentage** | 0.66 | 0.03 | 0.14 | 0.03 | 0.14 | |
| | | | | | | |
| **E** | 6 | 0 | 3 | 0 | 0 | 9 |
| K | 4 | 0 | 4 | 1 | 0 | 9 |
| B | 1 | 3 | 0 | 2 | 0 | 6 |
| **Total** | 11 | 3 | 7 | 3 | 0 | 24 |
| **Percentage** | 0.46 | 0.13 | 0.29 | 0.13 | 0 | |
| | | | | | | |
| **J** | 36 | 0 | 0 | 0 | 7 | 43 |
| L | 48 | 1 | 6 | 0 | 0 | 55 |
| O | 8 | 7 | 11 | 1 | 0 | 27 |
| I | 3 | 4 | 1 | 0 | 0 | 8 |
| **Total** | 95 | 12 | 18 | 1 | 7 | 133 |
| **Percentage** | 0.71 | 0.09 | 0.14 | 0.01 | 0.05 | |
| | | | | | | |
| **O** | 7 | 0 | 0 | 0 | 0 | 7 |
| I | 3 | 2 | 0 | 0 | 1 | 6 |
| **Total** | 10 | 2 | 0 | 0 | 1 | 13 |
| **Percentage** | .77 | 0.15 | 0 | 0 | 0.08 | |
| | | | | | | |
| **I** | 4 | 1 | 0 | 2 | 0 | 7 |
| **Total** | 4 | 1 | 0 | 2 | 0 | 7 |
| **Percentage** | 0.57 | 0.14 | 0 | 0.29 | 0 | |

**Figure 15**: Comparison of the ActivityTrailsLite data sets. The first participant listed is the one that recorded the data set.

The following six graphs compare participant's segments for the same data set. There are two for each device. The graphs are followed by a description of the graph and the pictures representing the event boundaries that at least two individuals had in common.

**Legend**

Event Boundary

Identical Event Boundary

Visual associated with the agreed segment

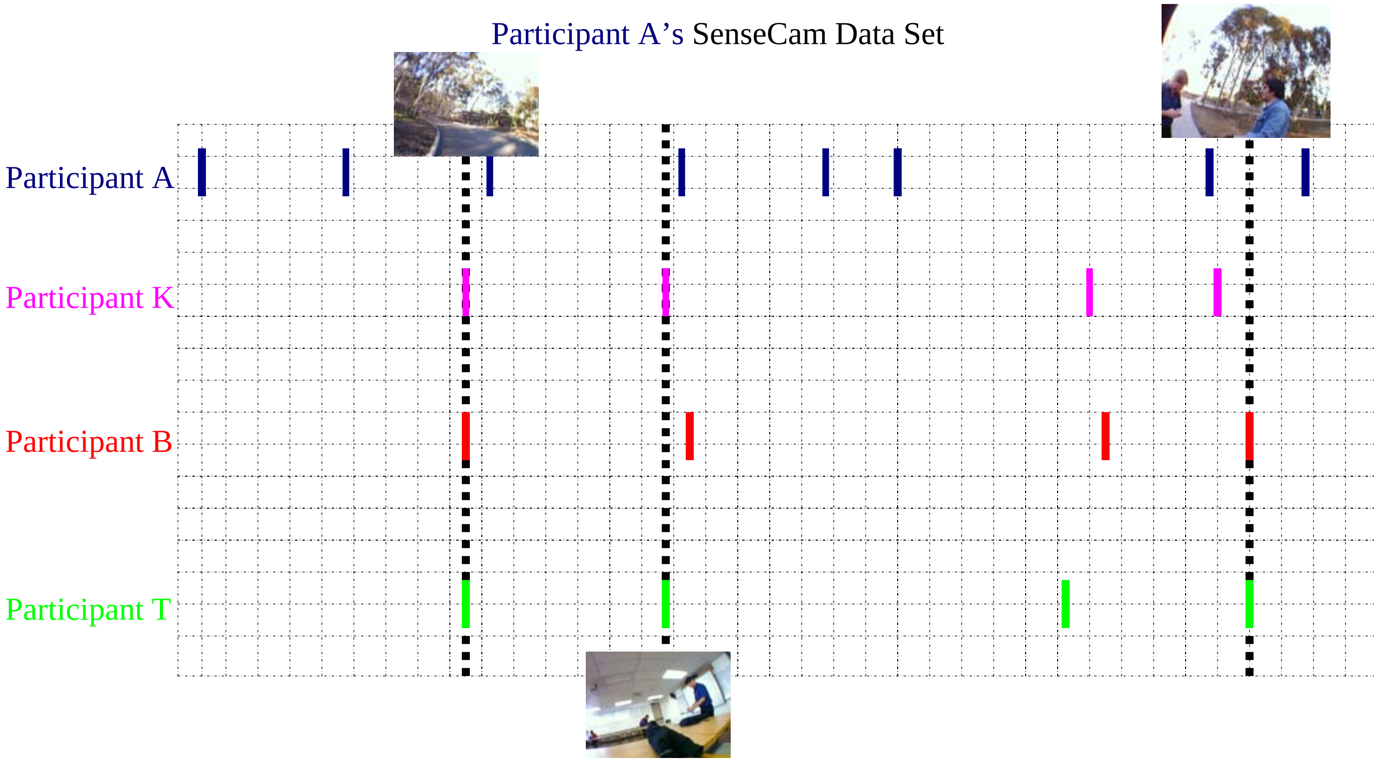


Participant A's data set illustrates discrepancy between segments between the four individuals. In this case the activity was fairly simple and straight forward, participant A left the career center and went to a presentation. Even in such a simple situation, there were still contextual clues that led to Participant A having more segments as well as discrepancies between the other four individuals, although they are very close in their segments and have the same number of segments. The pictures on the left and below are the points that at least two people agreed to put a segment.

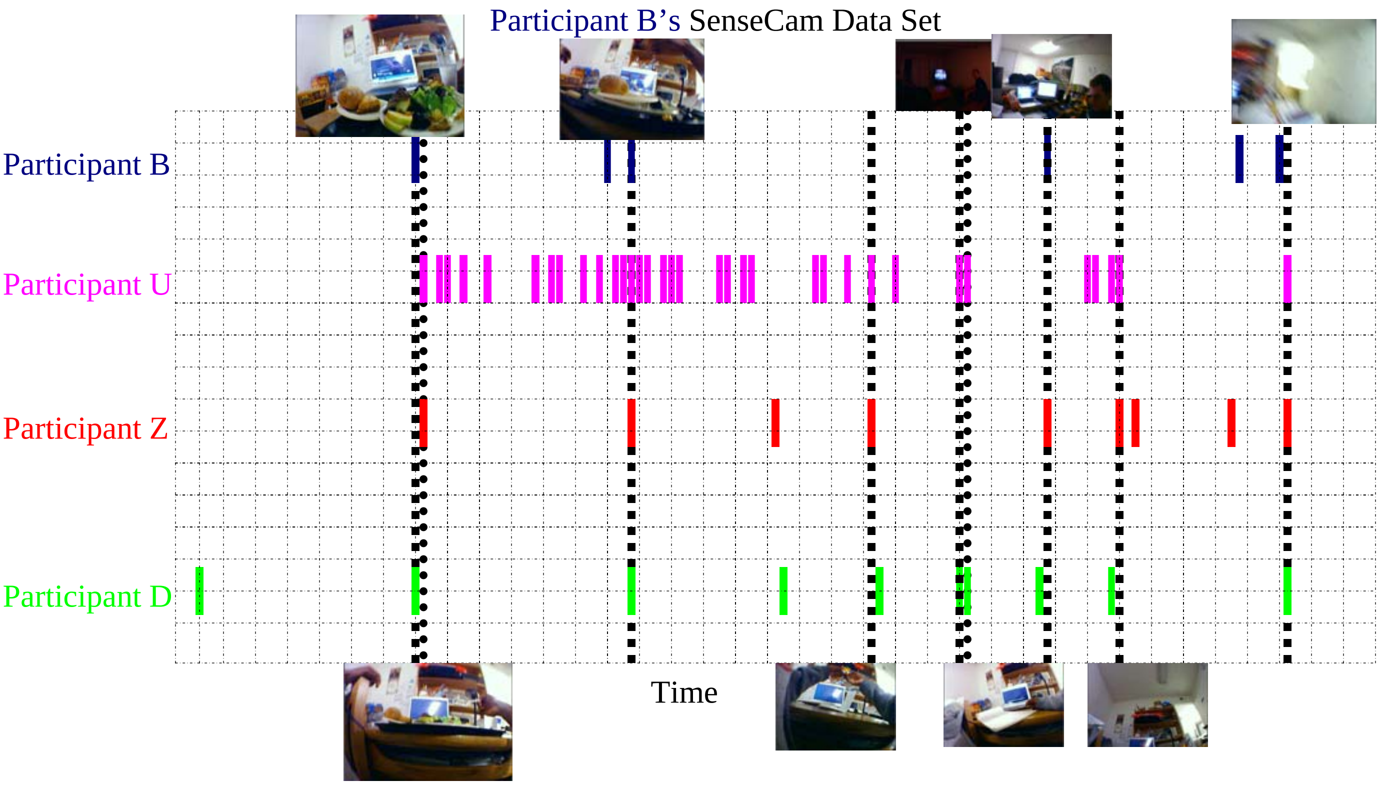


Participant B's data set, which was also collected using the SenseCam, was of a recorded evening at his home filled with food, videos, and homework. Because of the variety of tasks and individuals in the data set, the set was more complex and participant U had to utilize the Visual Cue segmenting strategy. Because of this, she has a lot of segments. The accuracy between individuals is also less in this instance than it was in the previous example, most likely due to the less structured type of activities occurring.

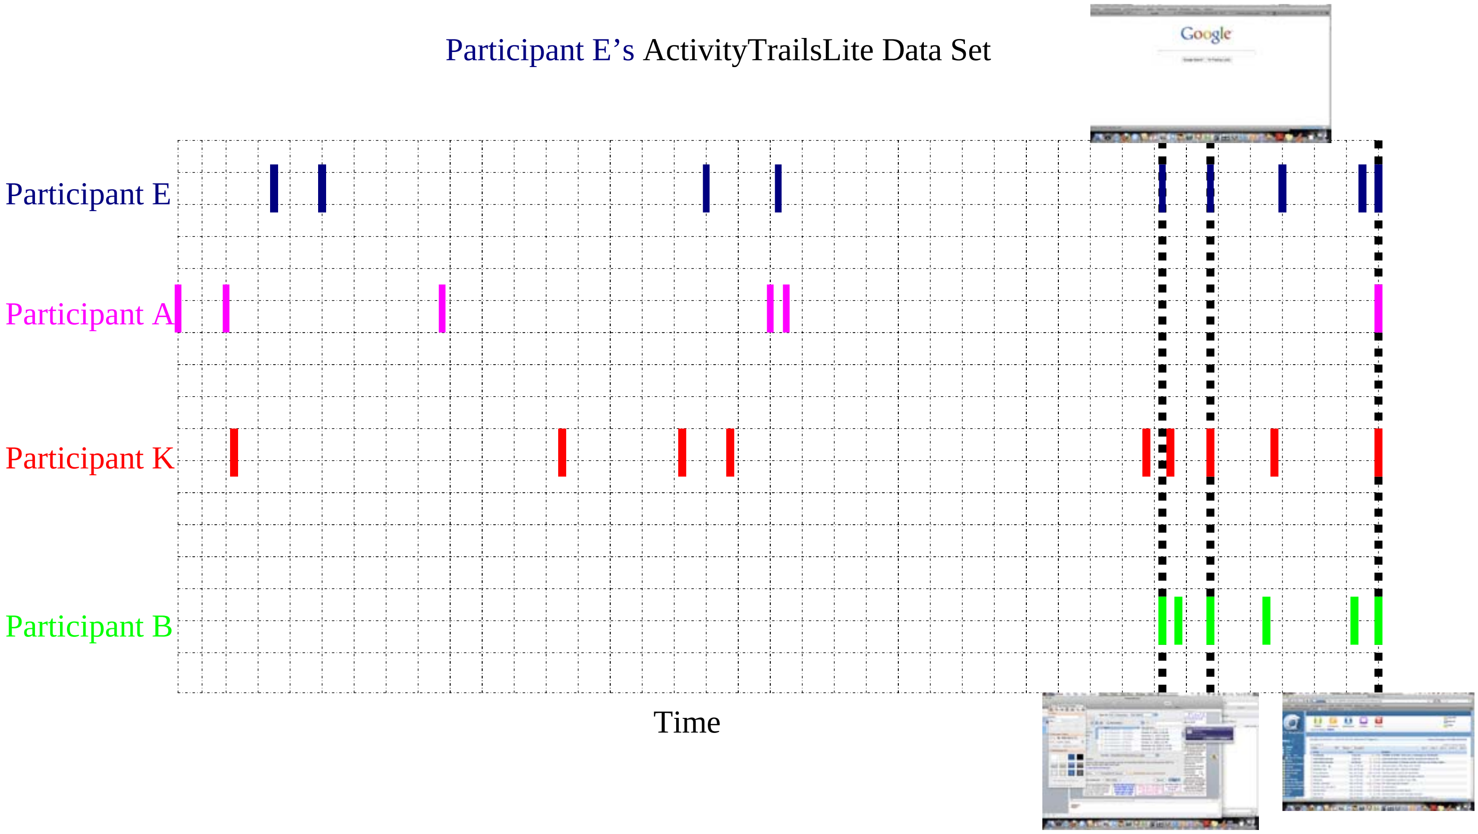


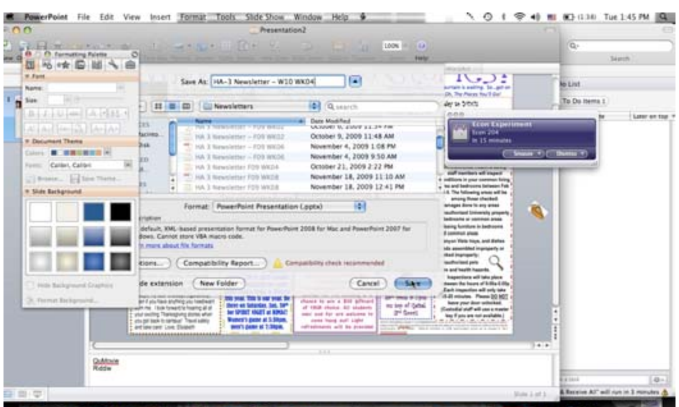

Participant E's data set recorded her making a resident advisor newsletter on her computer. The segments that agree are Natural Breaks, the first one she saved her document, the second one she left her document and went to google, and the third one she went to her email. These common points illustrate how these salient changes are seen as natural breaks in the task and have high agreement between participants.

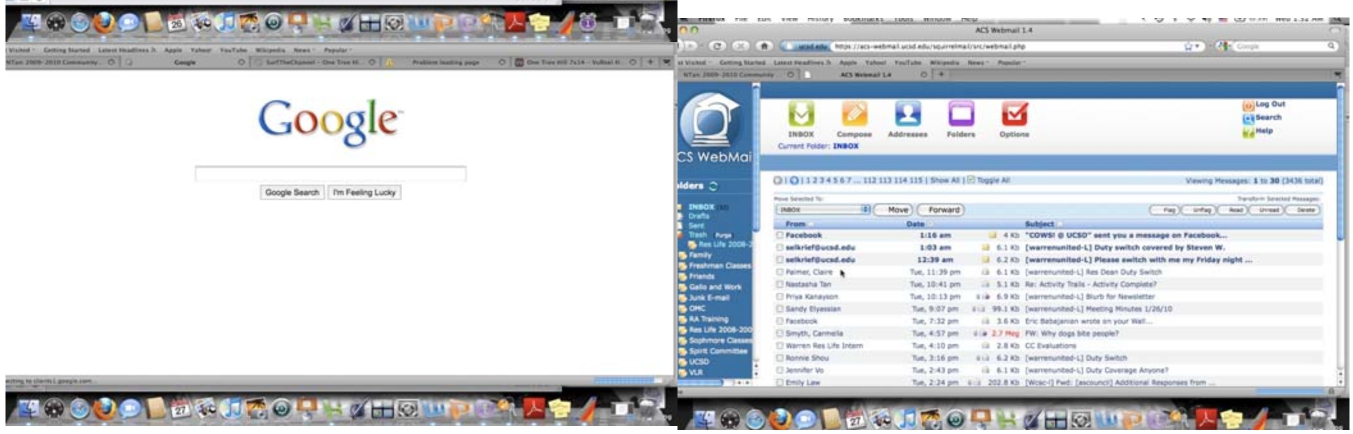

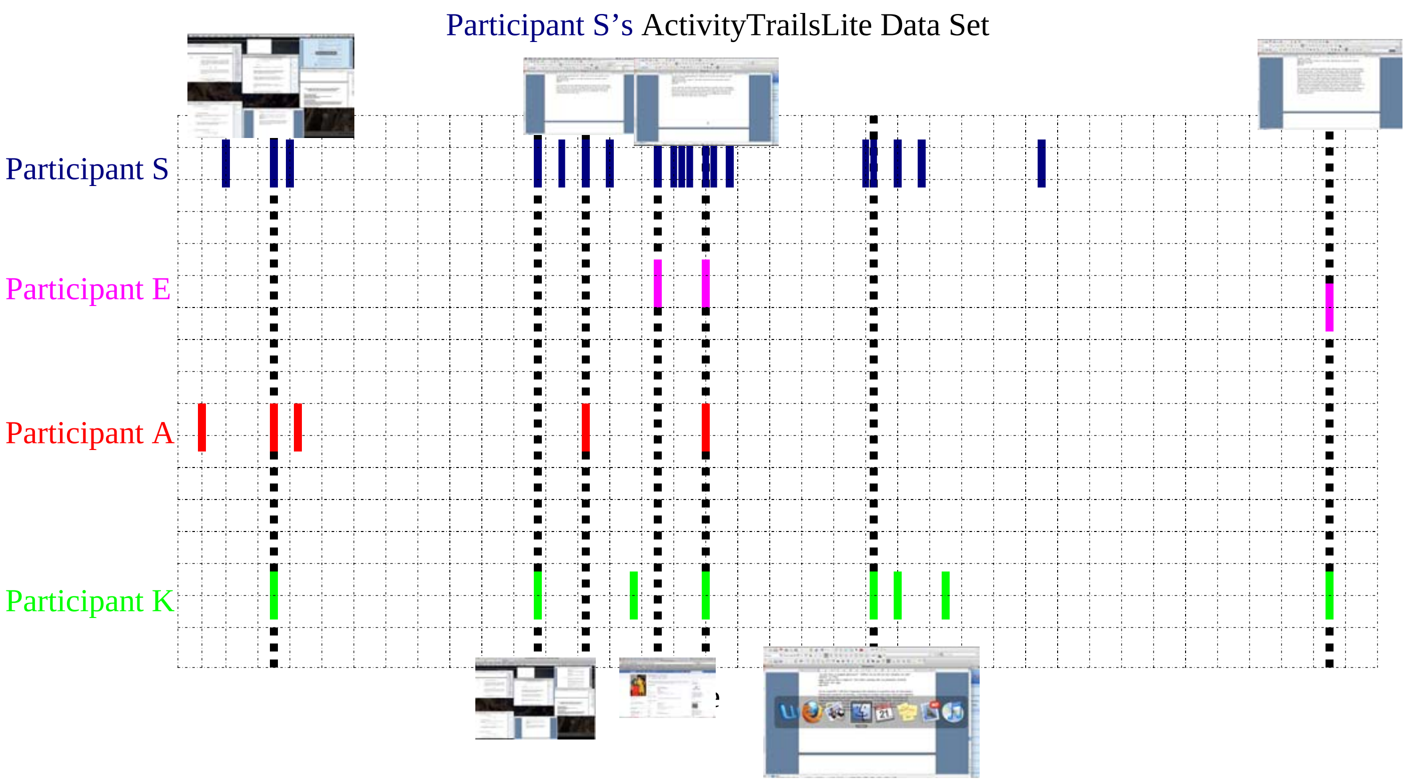


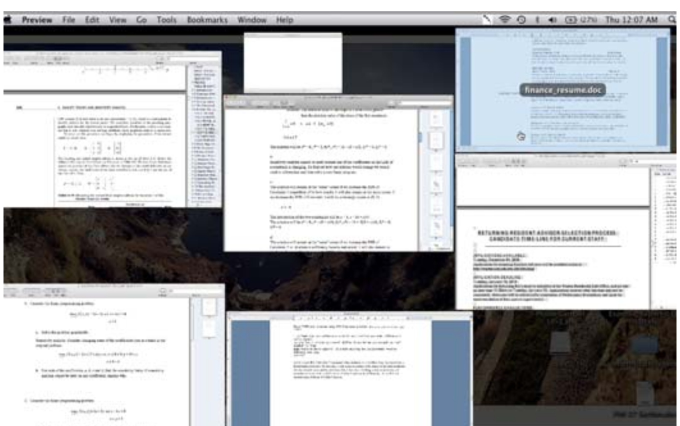

Participant S was filling out an application for a job during this data capture session. Participant S segmented half of his points because of contextual cues, making it fairly unlikely that anyone would be able to match his segments exactly. Because there was a lot happening not visible on the screen, he had many more segments than any other participant that segmented this data set.

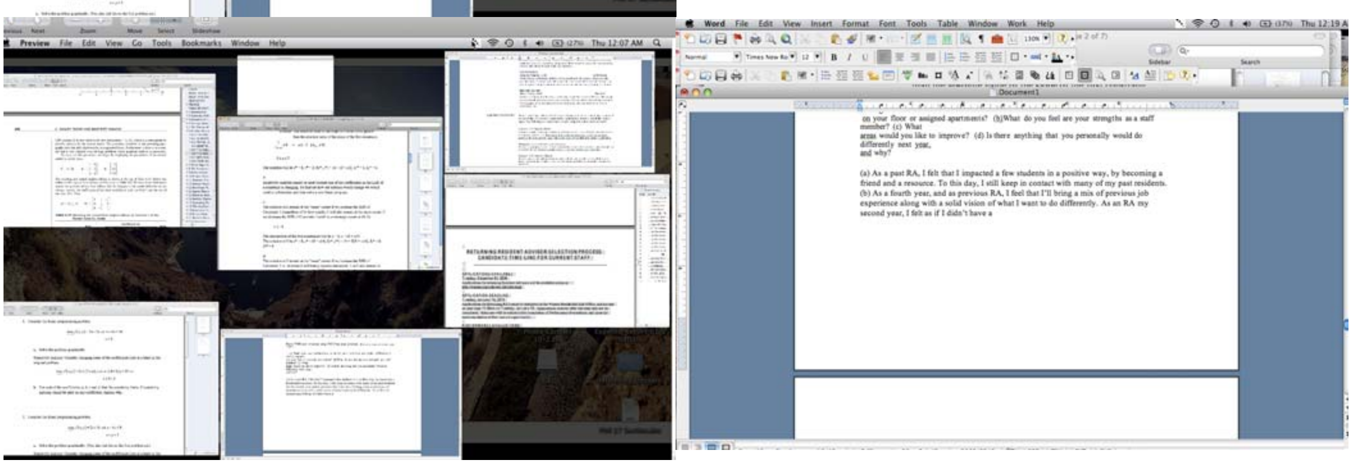

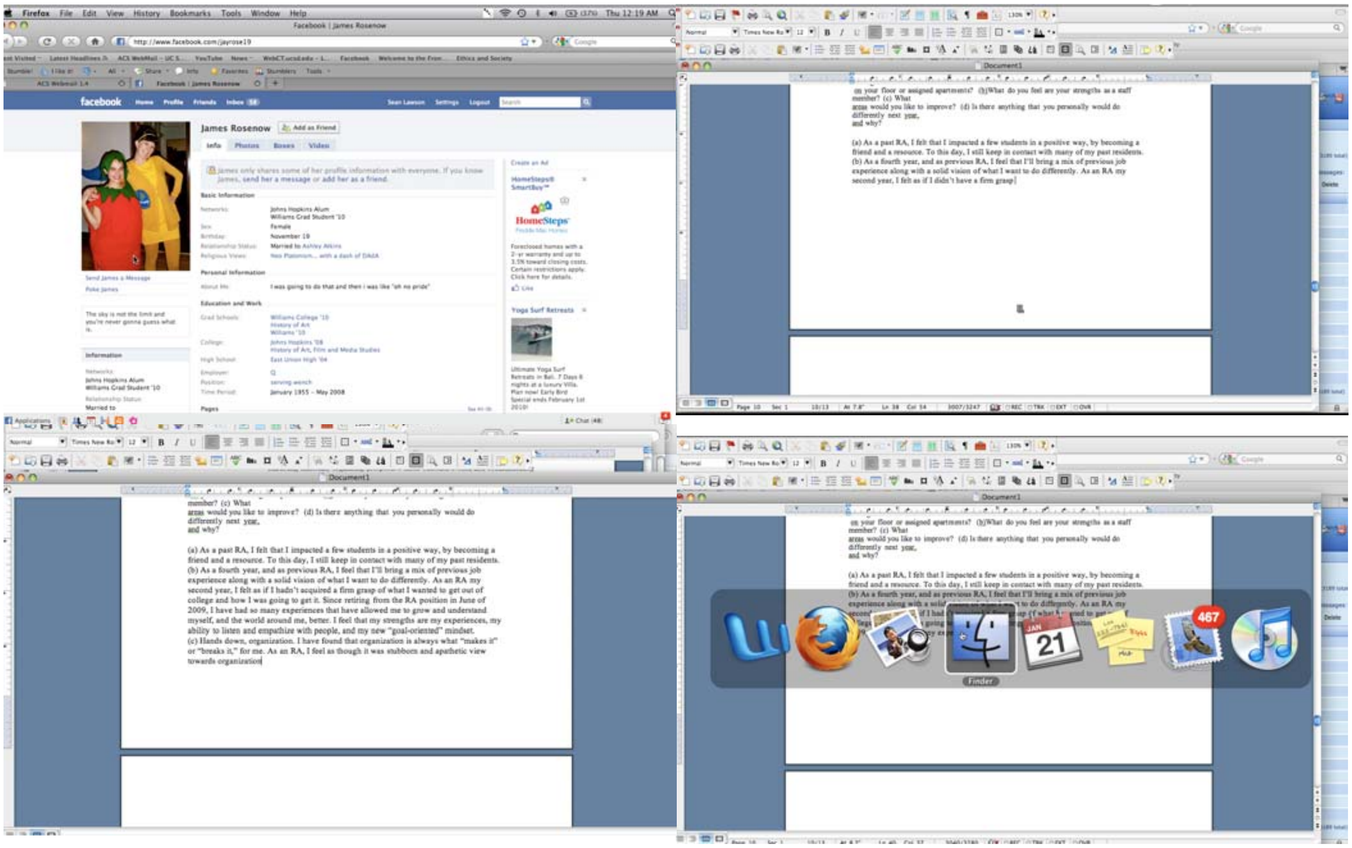

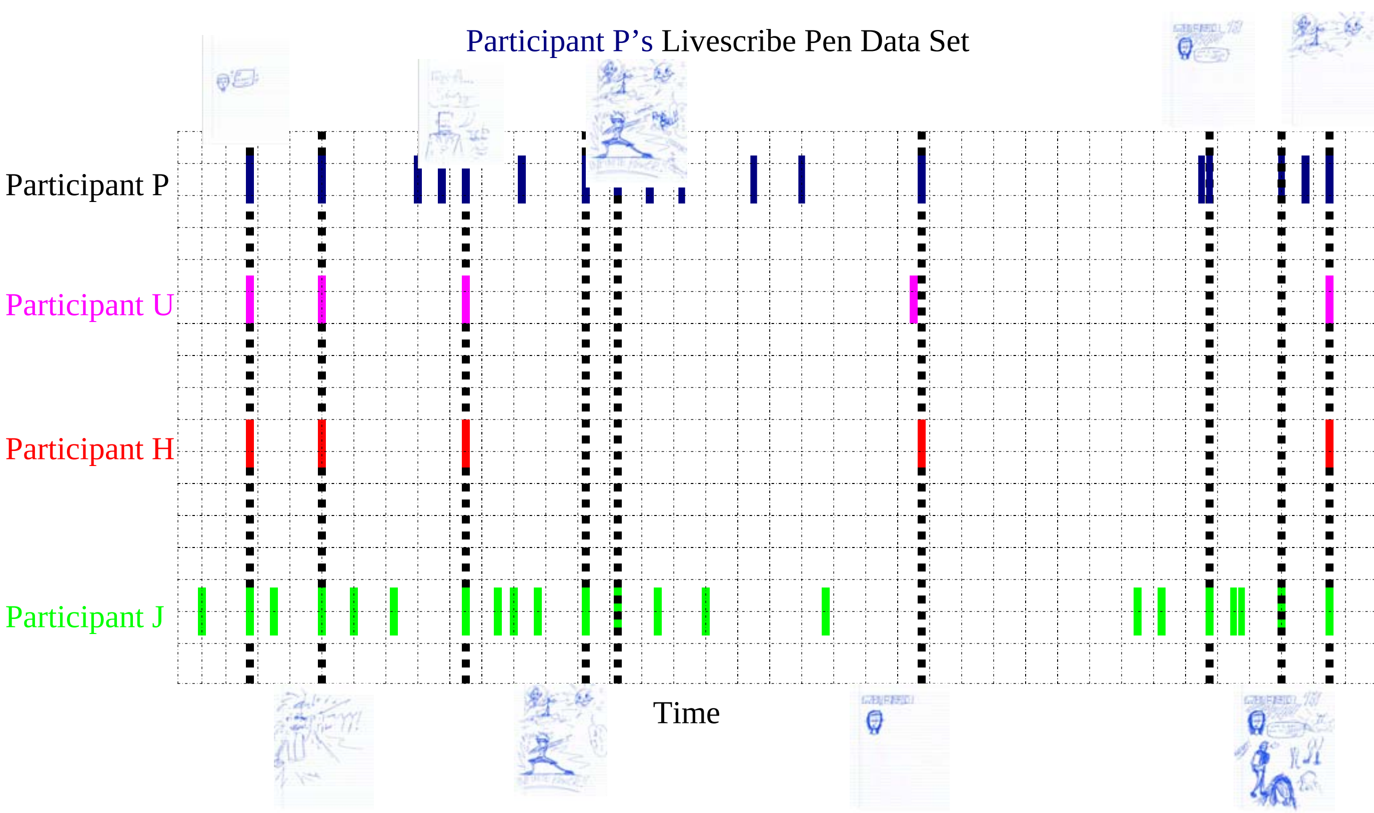
Participant P's Livescribe Pen Data Set
Participant P
Participant U
Participant H
Participant J
Time


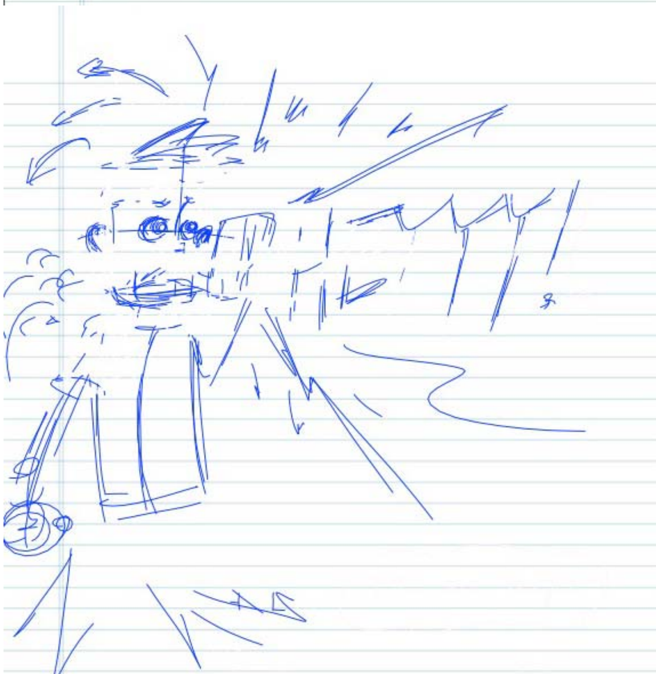

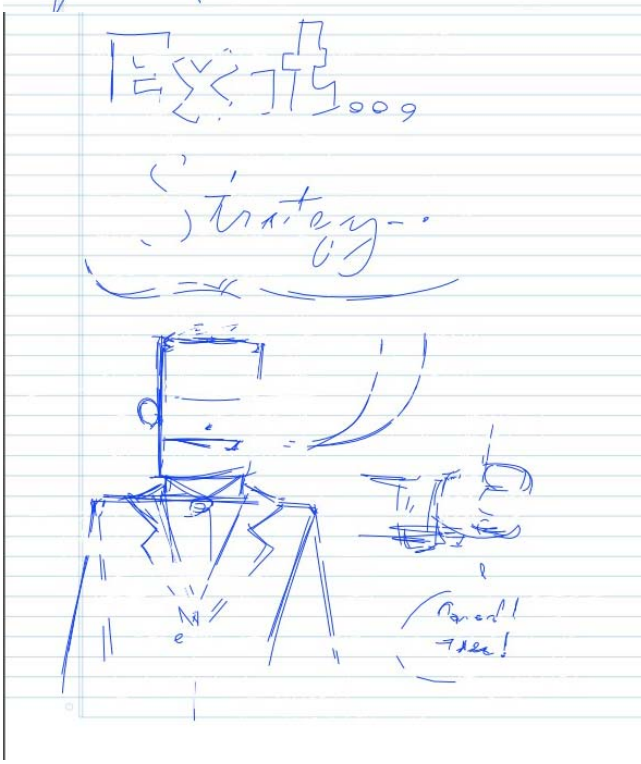


Participant P was inspired to draw most of these things based on the environment he was in and the conversations he was having. Many of his segments are Contextual Cues which are not visible to the other participants, especially in this medium. There are still quite a few segments where everyone agreed. This is because of the Natural Breaking point between each page. Everyone segmented the end of each


ROCK & ROLL!!
INFINITE POWER!!
by The Delfonics!


What if people think I'm CRAZY?!
SKETCHY!

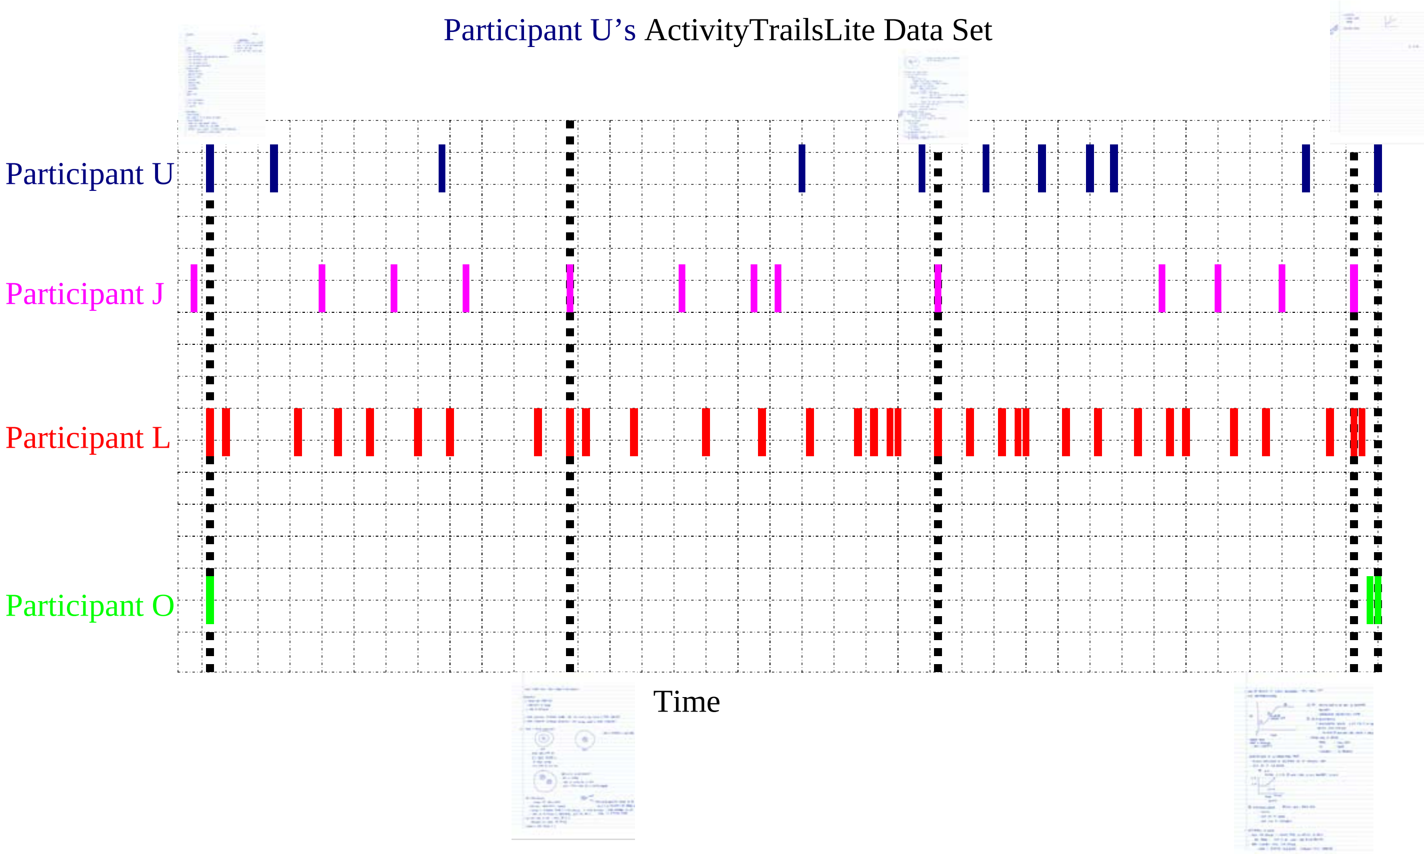



Participant U took notes during a lecture. In addition, she seemed a little uncomfortable with the auto-confrontation style of interviewing which may have affected her segmentation activities. The agreement here is again due to the Natural Breaks in the task, the end of each page. Participant O viewed this activity as pretty much one continuous activity, while the other three saw different points which to break up the task. This task was primarily segmented using Visual Cues because of the type of data available and the difficulty with coming up with a scenario to fit the task besides the general assumption that the individual was in a lecture.

- don't need IPTG - don't need to talk about it

- Antibiotics
  - 2 factors that affect ZOI:
    1. sensitivity of target
    2. rate of diffusion

1. same antibiotic - different targets - can you usually tell which is more sensitive?
2. same organism - different antibiotics - can you tell which is more effective?

(1) which is more sensitive?

rate of diffusion is the same

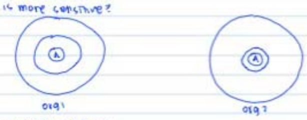


org 1 org 2

- more sensitive b/c
  [ ] near outside is
  ↓ than inside
- still able to kill org.


sensitivity to antibiotics?
- ZOI is same
- rate of diffusion is diff.
- don't know what ZOI is saying here


← resistance/sensitivity based on this
no [ ] of antibiotic @ certain p

ex: vancomycin
  10mm ZOI - resistance
- colistin - sensitivity (10mm)
  - which is diffusing faster? = vancomycin ← since resistant & same diameter as the other it's diffusing faster
  - rate of diffusion is completely diff for the 2
- for any ZOI in KB - 100% @ [ ]
  - desired kill rate (ie 95%)
- usually use single [ ]

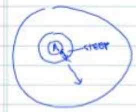


- KB method used when symptoms don't immediately tell you what the org is

- antibiotic fn : cellular targets
(1) cell wall synthesis inhibitor
  ex: penicillin
    - has β-lactam ring
    - prevents cross-linking of peptidoglycan
    - gram (+) & some gram (-) (narrow spectrum)
  - resistance - prdn of β-lactamase
  (possible) - prevent binding/transport
      - for gram (-) - LPS
  - modify drug - ampicillin - amino penicillin
      - helps agst LPS (cross LPS) → from narrow → moderate
    - amoxicillin - given as augmentin
      +
      clavulanic acid - gives amoxicillin protection agst β-lactamase
  - all 3 have β-lactam - (pen, amp, amox)
  methicillin - methyl added
    - modification of penicillin

(2) protein synthesis inhibitors
  ex: kanamycin (broad spectrum)
    - transposon can make it narrow
      - org. that can't transport can't be sensitive

own targets development of resistance

(3) nucleic acid synthesis
  - DNA synthesis
    ex: cipro - ciprofloxacin
  - RNA synthesis
    ex: rifampicin
(4) cell membrane disruption = lyse
  ex: polymyxin
(5) anti-metabolite - inhibits some metabolic pathway
  ex. sulfa-drugs (synthetic)

- look @ growth of vibrio natriegens - BHI - NaCl - 37°
- use spectrophotometer

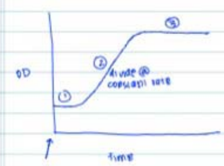


already have some in inoculum
∴ don't start @ 0

① lag - getting used to the env by producing enzymes
  - metabolically active
② log or exponential
  - exponential growth 2 → 4 → 8 → 16 →
  - growth rate constant
    - dividing @ maximal rate which is con
- change rate of growth
  - temp - toxic prdts
  - pH - space
  - nutrients - $O_2$ aeration

doubling time or g (generation time)
- always measured in log phase b/c of constant rate
- pick OD in log phase
  ex: 0.4
  double → 0.8 @ what time is this reached? 20 min

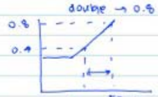


20 min

③ stationary phase   growth rate = death rate
  - toxins
  - run out of space
  - run out of nutrients

- variations in curve
  - long lag phase - taking time to adjust to env.
    ex: temp - 37° → 4° won't see much growth
  - goes straight into log phase
    - didn't change medium - already fully adapted

assignment due Thurs.!

- log phase
  - steep slope
  - height

- stationary phase

te or not

# Conclusion

This preliminary study has accumulated a massive amount of data that can be further analyzed to find many more relationships between the event boundaries, the device the data was collected with, and the comparisons between those and within those two variables. From this study, it appears that there is large variability in most aspects of event segmentation "in the wild." It appears that the most similarity between segments occurs when there is a Natural Breaking point. These were the most common in the Livescribe data because of the natural ending of a page. Participants who were at, or recognized, events that were captured were also able to use contextual cues in many cases to produce segments, even when they did not record the data set themselves. They were able to reinstate the context of the event in order to better break it up into meaningful chunks.

There was not a clear answer whether individuals see their activities as more of a continuous stream or divided into many interruptions. In most cases the participant who recorded the activity fell in the middle with regard to the number of segments for their data set.

Because of the number of variables and the amount of information, a more effective means to visualize it needs to be found, such as the techniques discussed by Edward Tufte [11, 12]. A multidimensional figure might be able to shed more light on this complicated activity to better help researchers understand how people segment and if there are any patterns that are currently hiding in the data.

In addition, there could be a trial round that determined whether people were comfortable with auto confrontation to help ensure that all participants are able to contribute fully to the study. There could also be more than one coded to help ensure the level of accuracy of categorizing each type of segment. There is a lot of future work that can be done in this area and hopefully this preliminary study has provided some insight into areas in need of further exploration.

## References


[1] C. Gurrin, A. F. Smeaton, D. Byrne, N. O'Hare, G. J. Jones, and N. O'Connor. An examination of a large visual lifelog. In *AIRS 2008 - Asia Information Retrieval Symposium*, 2008.

[2] B. A. Nardi, D. J. Schiano, M. Gumbrecht, and L. Swartz. Why we blog. *Commun. ACM*, 47(12):41–46, 2004.

[3] "SenseCam Wiki." Dublin City University. <http://www.clarity-centre.org/sensecamwiki/index.php/Welcome>

[4] C. Chen, *Information Visualisation and Virtual Environments.* London: Springer-Verlag, 1999.

[5] J.M. Zacks and K.M. Swallow, Event segmentation, *Current Directions in Psychological Science* **16** (2007), pp. 80–84.

[6] DCog-HCI Lab. University of Califronia, San Diego.

[7] D. Byrne, A. R. Doherty, G. J. Jones, A. Smeaton, S. Kumpulainen, and K. Jarvelin. The sensecam as a tool for task observation. In People and Computers XXI Interaction, 2008

[8] Doherty, Aiden R. and Smeaton, Alan F. "DCU SenseCame Browser." Dublin City University.

[9] "Introduction to SenseCam." Microsft Research. 2007 Microsoft Corporation. <http://research.microsoft.com/en-us/um/cambridge/projects/sensecam/>

[10] C. Hannon. Paper-Based Computing. *Viewpoint*. EDUCAUSE QUARTERLY, Number 4, 2008.

[11] Tufte, Edward R., *The visual display of quantitative information / Edward R. Tufte* 2nd edition Graphics Press, Cheshire, Conn. (Box 430, Cheshire 06410) : 2001

[12] Tufte, Edward R., *Envisioning Information / Edward R. Tufte* Graphics Press, Cheshire, Conn. (Box 430, Cheshire 06410) : 1990

[13] Doherty, A.R., Byrne, D., Smeaton, A.F., Jones, G.J.F., Hughes, M.: Investigating Keyframe Selection Methods in the Novel Domain of Passively Captured Visual Lifelogs. In: Proc. of the ACM CIVR 2008, Niagara Falls, Canada, (2008).

[14] A. R. Doherty and A. F. Smeaton. Automatically segmenting lifelog data into events. In *WIAMIS 2008 -9th International Workshop on Image Analysis for Multimedia Interactive Services*, 2008.

[15] Zacks JM, et al. Event understanding and memory in healthy aging and dementia of the Alzheimer type. *Psychol. Aging.* 2006;21:466–482.

[16] Newtson, D., & Engquist, G. (1976). The perceptual organization of ongoing behavior. Journal of Experimental Social Psychology, 12 , 436-450.

[17] Czerwinski, M. & Horvitz, E. (2002). Memory for Daily Computing Events. In Faulkner, X., Finlay, J. & Detienne, F. (Eds.), *People and Computers XVI, Proceedings of HCI 2002*, 230-245.

[18] Rekimoto, J. (1999). Time-Machine Computing: A Time-centric Approach for the Information Environment. In Proceedings of Annual ACM *Symposium on User Interface Software and Technology*, *UIST '99*, 45-54.

[19] González, V.M. and Mark, G. "Constant, constant, multi-tasking craziness": managing multiple working spheres. In *Proc. CHI'04*, ACM Press (2004), 113-120

[20] Cohen, C. E., & Ebbesen, E. B. (1979). Observational goals & schema activation: A theoretical framework for behavior prediction. *Journal of Experimental Social Psychology, 15,* 305-329.

[21] Hodges, S., Williams, L., Berry, E., Izadi, S., Srinivasan, J., Butler, A., Smyth, G., Kapur, N., and Wood, K. (2006). SenseCam: A retrospective memory aid. *Proc. Ubicomp 2006*.

[22] Sellen A. J., Fogg A., Hodges S., Wood K. Do Life- Logging Technologies Support Memory for the Past? An Experimental Study Using SenseCam. *In Press to appear in. CHI 2007 Conference on Computer Human Interaction, NY: ACM Press.*

[23] O'Connail, B. & Frohlich, D. (1995). Timespace in the workplace: Dealing with interruptions. *CHI '95 Conference on Human Factors in Computing Systems,* Extended Abstracts, ACM Press, 262-263.

[24] Zacks J.M., Speer N.K., Vettel J.M., Jacoby L.L.*(*2006*).* Event understanding and memory in healthy aging and dementia of the Alzheimer type*. Psychology & Aging, 21, 466–482*

[25] Aizawa, K. Hori, T. Kawasaki, S. Ishikawa T (2004). Capture and efficient retrieval of life log. In: *Pervasive 2004 workshop on memory and sharing of experiences*

[26] Song, H. Guimbretiere, F. (2009). The ModelCraft framework: Capturing freehand annotations and edits to facilitate the 3D model design process using a digital pen. *ACM Transactions on Computer-Himan Interaction (TOCHI) Volume 16, Issue 3*

[27] Qi, J. and Buechley, L. Electronic popables: Exploring paper-based computing through an interactive pop-up book. *Proceedings of TEI '10*, pp. 121–128, NY: ACM, 2010.

[28] B. Erol, E. Antunez, and J.J. Hull, "Hot-Paper: Multimedia Interaction with Paper Using Mobile Phones," *Proc. 16th ACM Int'l Conf. Multimedia*, ACM Press, 2008, pp. 399–408.

[29] Mollo V, Falzon P (2004) Auto and allo-confrontation as tools for reflective activities. *Appl Ergon 35(6):*531–540