Table of Contents

Background and Motivation

The Internet is a tool used for an extensive range of purposes. Some very common uses include searching for information, enjoying multimedia entertainment, and interacting with other people. Many steps are involved in order to achieve these end-goals, and occasionally some part of these steps are unable to be satisfactorily completed, which inhibits achieving the desired result. These inhibiting interruptions are commonly referred to as “breakdowns” within the workflow of a task. Because all of the members of our group are frequent Internet users, we have all experienced the consequences of encountering a barrier in the form of a breakdown which prevents us from completing a task.

As a result of our collective experiences when using the Internet, our group decided to examine the way that tasks are completed within the framework of the browser, specifically with regard to browser tabs. The original goal of our project was to redesign browser tabs to provide improved context-dependent indications about a user’s motivation behind each specific tab’s creation. We decided to gather data from experienced Internet users – users who already had developed a sort of ‘problem blindness’ to these breakdowns and had strategies in place for dealing with nuances. Our original data indicated that these users experienced the most breakdowns while completing specific patterns of activity. Additionally, these users experienced less negative impact from lack of context than we had initially anticipated. The only activity which permeated our entire data set was the act of searching for information, specifically using the Google search engine to seek various types of knowledge. After enough data had been gathered to indicate the importance of this task, we collectively decided to focus

our design efforts on changes to the browsing experience which facilitate the ability to obtain preferred search results quicker, and manage those results most efficiently.

Many factors were taken into account when our group decided which Internet user demographic we would be observing. We chose to observe college-age students, not only because of their immediate availability to our group members, but mainly because of their proficiency with the Internet. Additionally, because our initial project scope involved redesigning browser tabs to provide improved context indications for any type of browsing activity (information search, leisure, socializing, etc.), our team felt that the college-age student demographic had the most familiarity with all of these task categories. When we reanalyzed our project's scope and decided to focus on facilitation of information search, our target users proved to be experienced at the task and provided a rich dataset of activity information.

As a group we utilized a user-centered design methodology to enable us to incorporate actual user practices into the design process. Also known as contextual design, this method of data retrieval, analysis, and ultimately incorporation encouraged us to gather our data through first hand contextual interviews conducted with real users of the Internet within their actual workspaces. In order to obtain data that best represented actual Internet use scenarios, our group members selected users who we were already personally familiar with to foster comfort within the user and to encourage natural habits which may have been inadvertently concealed among strangers.

Methods and Results

Because this was the first time we were introduced to the contextual design process, we tried to adhere as closely to the techniques described in the textbook as possible. However, given the short amount of time that we had to complete the project, we did have to evolve some of the methods to fit our needs. Overall, we approached this as a learning experience, and we were constantly receptive to feedback from the TAs and Professor Hollan. The great value we placed on following the process proved to be worth the extra effort because of the rich data we were able to extract.

Group communication was a strong point from the onset. Since we are all Gmail users, we relied on Google Docs, Google Calendar, and Gmail threads to share and coordinate information. Our meetings were a time to share, in person, everything that we had been working on individually in the previous week. Initially, these meeting were held on Mondays, then gradually progressed to Mondays and Wednesdays, and finally to every day during the end of week eight, week nine, and the beginning of week ten. Our typical meeting lasted anywhere from three to four hours, and were consistently four to five hours toward the end of the quarter. At the beginning of these meetings, we identified the goals we wanted to accomplish and took detailed notes of everything that was discussed [Binder Documentation VI]. At the end, we summarized and planned out our concrete goals for the coming week.

First Round Contextual Interviews

Methods

The first step that we took was identifying users to interview. We drafted an "Interview Request" letter that we sent out to "power-users" of the Internet. As mentioned previously, these were users that used the Internet on a daily basis, and were comfortable enough to have already developed some sort of problem blindness and strategies for dealing with nuances and breakdowns in their workflows. After we received responses and identified what we would be looking for in the interviews, we all set out to conduct our first contextual interviews. We decided to conduct the first round of contextual interviews with people that we knew due to the personal nature of the task. We asked these frequent Internet users not to clean up or close tabs on their desktops, as we took a note of which tabs were already open at the beginning of the interview. We asked that they schedule a time to meet with us when they had a task they wanted to work on such as researching, completing an assignment, or even searching for an item to purchase. Our goal was to observe them during a natural Internet browsing session so that we could identify real breakdowns that they experience on a daily basis.

At the beginning of the interviews some things we identified were how long the user has been using the Internet, which browser they normally use, and what they usually use the Internet for. Obtaining this background information provided us with a solid building block on which we could analyze their actions and generate questions. Our interviews lasted anywhere from 45 – 90 minutes; we found that this amount of time allowed the user to overcome any initial awkwardness or discomfort they may have at first experienced. We practiced the "Master-Apprentice" model of interviewing in which our group member acted as an 'apprentice' learning the task at hand from the user. We took caution to always ask clarifying questions if their intent for any action was not completely obvious. Some things that

we specifically kept in mind and made an effort to capture in our notes were: how users access new content (tabs versus windows), spatial organization (how they mark what is important), keyboard shortcuts, and mouse shortcuts. Overall, these points can be summarized into *how users structure information in their mind and how they translate that into browser tab usage.* In addition to the many pages of handwritten notes that we acquired [Binder Documentation I], we also made audio recordings of the interviews, which we used to supplement our notes later in the quarter prior to our data consolidation process.

Upon completion of our first round of contextual interviews we reflected on the process as a group and discovered a few issues. First, making the user feel comfortable proved to be very important and not always easy. Our interviews ran on the basis of a natural browsing experience, and if the user did not feel comfortable with us observing their actions and taking notes then our data would not be usable. We found that approximately the first 15 percent of the interview was the most difficult, but after this initial period, most of the users began to embrace the process. Secondly, we found that taking notes was sometimes very difficult. Sometimes we missed actions that the users took, prompting us to ask them to retrace their steps. Additionally, asking many clarifying questions occasionally interrupted their workflow, but did prove to be extremely helpful for identifying their motivations behind their actions. Lastly, we found that not all questions led to useful information; it became very important to enforce the key concepts we were looking for and avoid falling into “ratholes.”

Results

Due to the difficulty of scheduling meetings with such a large group, we decided to hold our interpretation sessions individually or in pairs, depending on the availability of team members. Because this strategy minimized the elapsed time between the conclusion of the interview and the interpretation session, we adopted it over a more discussion involved interpretation session that would inevitably be more difficult to schedule and cause details from the interview to be lost.

As soon as possible after an interview, the interviewer (and in some cases a partner) reviewed his or her notes and reflected on the interview. During this time, the interviewer annotated interesting sections, common patterns, noticeable breakdowns, and even design ideas. Everything that was important or relevant was recorded in the interviewer's notes so that nothing would be lost between the interpretation and the group sharing sessions.

Our first interview subject, L1a, was an interesting introduction to the contextual interview process for our team. For this interview, we tested out audio recording as a method to complement paper note taking. The subject was the only one of our users that used the browser "Opera." Seeing the features of Opera that this user appreciated was our first clue into the influence visibility had on the browsing experience. During the interview, he also took the time to show us the way he used "Internet Explorer," and he was one of the only subjects that said that he switched between browsers frequently. Because this was our first interview, it was also an excellent learning opportunity for the entire team in terms of what strategies were effective during the interview for extracting interesting data, and what strategies were not effective. The interview team found that asking open-ended questions and asking pointed questions about behaviors that seemed interesting or unique was the most effective way to get

the user to talk about his browsing habits without disrupting his workflow too much. This user started a search in a fresh tab and open links he was interested in in consecutive new tabs. He then visited each of the open tabs in order, closing tabs that he finished extracting information from, or that were not relevant, and sometimes left tabs open if he thought that he would reference them in the future. However, when the user found a link to be particularly interesting, he viewed the newly opened page immediately after opening it in a new tab.

Our next interview subject, N1a, had a different approach to searching than L1a did. In some cases, after opening interesting links in consecutive tabs this user quickly analyzed the information on each relevant page and flip to the next if they were not able to quickly find what they were looking for. If none of the tabs that she opened from their initial search yielded the information that she wanted in an easily digestible format, this user either returned to the original search tab to further refine her search, opened a new tab and started a new search, or worked her way back through the information that she skipped over, this time reading more carefully (i.e., she worked harder to extract the relevant information).

D1a, our third interview subject, provided extremely interesting data because it was almost the exact opposite of another of our user's (M1a, discussed later on) data. When the interview started, this user already had 13 tabs open in the current browsing window, and when she viewed all her windows at once using the Mac exposé feature it became apparent that she had multiple windows open with the same number of tabs in each. When questioned about this behavior, this user replied that she really disliked closing tabs because she feared running into the problem of needing to reference them later. She instead preferred to open as many tabs as could fit into the browser window. When the window became so full that she

could no longer navigate efficiently, she just opened a new window and started the process over again.

Our fourth interview subject, L1b exhibited similar patterns to those of L1a, in that he worked his way through lists and opened the interesting or relevant links in new tabs. In L1b's case, the user was searching for new music that he wanted to add to his phone and/ or computer playlist. In contrast with L1a's interview, the interviewer asked more questions about L1b's thought process as he opened new tabs, switched to a different website, etc. It was interesting to note in L1b's actions that when he scrolled through pages of songs on YouTube, his fluid system of simply opening a handful of new tabs from links was disrupted. Since YouTube videos start playing instantly when they are opened, L1b had to go to the page of each newly opened tab and pause the video before returning to the list of videos to continue scrolling through songs.

Our fourth interview subject, Co1a, exhibited some previously unseen behaviors. This user openly admitted to preferring to have fewer tabs open at one time. His reasoning for this seemed to be his unfamiliarity with the capabilities of his computer. At different points during the interview, he expressed his concern that having too many tabs or windows open at one time might "overload" his computer.

Our next user, Cr1a, showed a similar pattern to those seen in the interviews with L1a and L1b. While trying to find information on how to build a mobile app., Cr1a opened links that he thought would be relevant into new tabs, and clicked on each one after having a handful open. He closed tabs that were irrelevant, long, or that he had previously seen, and he only liked to keep open tabs that seemed important to him.

User M1a was different from all the previous users, in that he used his tablet rather than a desktop or laptop. While he liked to “save” tabs (similar to user Cr1a), rather than accessing page content at a later time, M1a’s intention for keeping certain tabs open was to compose a tab list of pages that he wanted to open on his computer at home. In addition, M1a tended to reuse tabs (typing a new search in an already-used URL bar or search bar). M1a also liked to divide tasks by window, but never liked having more than a few tabs open in each window (if a task had more than a few tabs open, he opened a new window to “continue” the task).

Like some of the users described above, A1a also opened links in new tabs when she wanted to find out more about the content on a given link’s page. In addition, A1a opened a new tab for each of her activities (looking for a restaurant, checking Facebook, and checking her online banking). The most interesting note about this user was that she opened new tabs for new tasks rather than reusing them, even when she was done with an initial activity. After completing each activity, the user closed the corresponding tab.

Sharing Sessions

After our team members completed their first contextual interviews, our team met to begin to analyze our collected data, and to extract breakdowns that were found among multiple users. During this time we set out to describe our interview data to the rest of our group members, point out what was noted as particularly (based on the interpretation sessions), and examine any issues the users encountered in order to identify potential inspiration for design ideas. Dylan wrote notes on a chalkboard to organize our user data

visually, while Crystal utilized Google Docs to provide the entire group with notes to reference later in the design process [Binder Documentation VI].

One by one, our group members introduced their user by assigning them a code name, and walking the rest of the group through their contextual interview chronologically. When discussing the interview data, special attention was paid to what action the user performed and how they chose to pursue their goal. Focus was also placed on interesting data points which were identified in the interpretation sessions such as task methodology or information management, especially when these had been witnessed within another user's contextual interview. Additional focus was given to unique user traits which may be considered outliers. We took special care to thoroughly interpret user data before determining that a behavior or mentality was not representative of a problem relevant to the scope of our project, or not indicative of a breakdown of a larger group of users.

After our group examined each member's contextual interview data and noted important issues on the chalkboard, we analyzed our findings as a team. We drew lines from the issue to what we believed was the overall cause of the issue, and then to the root of that cause (e.g., Lots of tabs → Organization? → Visibility?). When we believed that one issue had multiple causes, we used lines to connect relevant cause-effect pairings. Ultimately we were left with a very descriptive diagram of how our users' detail-oriented "low-level" problems related to the goal-oriented "high-level" catalyst which spawned them.

Once we identified where the problems we observed our users experiencing had originated from, and were able to consolidate seemingly distinct low-level problems into related high-level groupings, we wrote the question that we believed we were now working

toward answering on the chalkboard. At this point in the project, that question had evolved to, "How do users organize information in their minds and thought processes, and by extension how does that translate into the way they use the browser?" We then noted what we believed to be an important concept while progressing with our project: Human-centered browser windows. Finally, below these concepts we listed some design considerations which we planned on further examining through the contextual design process. These concepts included Switch Costs, Context/Organization/Categorization, Screen Real Estate, and Applications with relation to the browser. All of these design considerations stemmed from the user issues which were most pervasive in our initial data.

Affinity Notes

As a group, we found it extremely helpful to generate our affinity notes during our individual (or partnered) interpretation sessions. Being new to the contextual design process, we found it difficult at times to separate the important information from the seemingly relevant but distractingly peripheral information. However, the original affinity notes that we did manage to generate provided us with a basis for consolidating our broad data set, and in general they provided us with enough of a common grounding to see the patterns in our data that led us to narrow our focus during our second round of interviews.

Although we were unable to generate an affinity wall with the number of affinity notes we generated from this first round, we did find it very helpful to create a pseudo affinity diagram. During a group meeting, we all translated our affinity notes to stickies. The first member of the group read an affinity note aloud and then stuck it to the oversized paper hung

on the wall. Then, taking turns, the group continued this process. One member at a time read the content of their sticky note out loud, and as a group, we decided if it could be clustered with another group of stickies or if it belonged on its own. We continued this process until all of the sticky notes had found a place on the wall. Then we did a short wall walking session. Instead of moving stickies individually as the book suggested, we chose to consult the group when we felt that one sticky belonged more strongly to another cluster. This method proved effective for our purposes.

Sequence Models

Because our first round of interviews were so loosely structured, our users completed a wide variety of tasks during their interviews. Such tasks ranged from starting research papers and searching for internships, to referencing methods of programming and casually browsing sites like Reddit. Originally, we planned to team up to create our sequence models. After giving that a try, we realized that the vastly different tasks and individual styles of note taking made partnering up a waste of time. Instead, we each completed sequence models for the interviews that we conducted.

For this round of interviews, we opted to complete the sequence models using a Google Doc spreadsheet. Naively, we thought that this would provide us with an easy way to make comparisons across tasks. Each group member inputted information about the task, the steps involved with the task, and tried to also include information on the user's intents, triggers for certain behavior, and anywhere that breakdowns occurred.

This process had some merits—it provided our group with some practice in thinking through the step by step processes that our users took to accomplish certain tasks, which led to both more affinity notes and the observation of subtle patterns that we missed in our initial interpretation of the interviews. However, probably the biggest benefit of completing these early sequence models was that it became glaringly obvious how widely distributed our data were and how difficult it was to consolidate data from such widely varied tasks. Because our users completed such incompatible tasks, comparing strategies was nearly impossible. In this way, although the process was frustrating to complete, building our sequence models heavily influenced our group strategy moving forward.

Data-Driven Extractions

Based on the affinity notes and sequence models, we were able to identify four major patterns/models in our data, and designated a unique name for each pattern: Curation and Consumption, Path of Least Resistance, Anxious Joe, and ADD Tom.

Curation and Consumption was a pattern seen in interviews with L1a, L1b, D1a, Cr1a, M1a, and A1a. These users began as "curators" by producing visual lists (opening links in new tabs) that indicated pages that they wanted to visit. After curating the lists, the users clicked on each new tab (in the order that they were opened), and decided whether they wanted to keep it open or to close it based on how relevant the information was to what they wanted. We called this behavior "consumption." Sometimes the users scrolled through a single page of search results at a time before looking through the tabs, and other times users seemed to transition from the curation process to the consumption process after a handful of tabs

waopened. However as seen with L1a, a user might cause a break in the curation and consumption pattern when a link's content is outstanding in some way. In such cases, a user might choose to view the page of the newly opened tab immediately after opening it in a new tab, rather than going through the tabs in the order they were opened.

A model that we identified is Path of Least Resistance (N1a emulated this pattern most closely). When this user was on a content-heavy page that she thought what she was looking for would be, she tended to avoid working her way through the content. Instead, she started a new tab to search for what she was looking for in a more direct way. Whenever a new idea (that the user thought would be another shortcut to the answer) sparked, the user immediately opened a new tab to pursue the new idea. If the new idea turned out to be irrelevant, the user abandoned the idea and returned to the previous tab. When and if the user no longer felt that she could not find the answer quickly, she returned to her original "safety net" and sifted through the content until she found the answer(s) she needed.

One model that we extracted from our data is characterized as Anxious Joe. One of the technical concerns that Co1a showed signs of distress when he had too many tabs open because he was worried his computer would become overloaded. For this user, we found that due to his relatively limited understanding of the browser's capacities, he was limited in the way that he searched the web. The Curation and Consumption pattern does not apply to Anxious Joe because he fears that his computer will not be able to handle the load. Thus, Co1a closed tabs immediately after using them, and in cases that he had to refer to previous pages, he clicked on the "back" button until he found the page he wanted. Because of his ways of navigation, Anxious Joe was extremely limited in how he searched and accessed information.

We characterized the final model that we observed as ADD Tom. M1a was a very tech-savvy and organized interviewee. However, he was so concerned with how his tabs were organized that he seemed a bit irrational. He equated productivity with never having more than a few tabs per window open (where all tabs pertained to the same task or subject), which resulted in him having a large quantity of windows open. While one may argue that his desire to divide tasks by windows was reasonable, it is important to emphasize the fact that M1a told the interviewer that even if one task required more than a few tabs, he would still open more windows for the same task.

Refocusing

After finishing the consolidation of all of our first round data and developing our major patterns, we realized that our results could not possibly be an accurate summary of the entire user population. We postulated that they probably were not even an accurate summary of the subset that we managed to cover. This revelation, combined with the fact that our very broad data set proved nearly impossible to consolidate, led us to the conclusion that we needed to focus in on an interesting subset of our findings, then go out and gather more data about that subset. We needed to accurately identify patterns in our data in order to generate design ideas that could have a meaningful impact. In order to identify patterns, we needed to extract linear data that we could then compare across different users.

In our analysis of our first round of interviews, we identified two behaviors of users (Curation and Consumption and Path of Least Resistance), and two user types (Anxious Joe

and ADD Tom). While the user types were somewhat interesting, they were not supported by an extensive data set. Intriguingly, both behaviors that we identified were deeply grounded in our thinly spread data. After much discussion our team decided to hone in on these behaviors. We noted that both behaviors were observed when users were in the process of actively trying to locate specific pieces of information. In their own ways, both behaviors were essentially strategies for efficiently accessing information and refining search.

These promising pieces of insight, centered around searching and search refinement, led our team to focus our research on strategies our users employed when they had specific pieces of information they were searching for within the browser. In order to generate more usable data centered around this specific focus, our group designed specific user tasks for our next round of contextual interviews. This would allow us to linearly compare different user strategies across the same range of tasks.

Second Round Interviews

Methods

For our second round of interviews, we decided to modify the contextual interview model and give our users two predefined tasks that they must complete in the allotted interview time. We observed different users than those in the first round of interviews, but still sought out the same type of power-users who we were familiar with in order to facilitate a comfortable environment. We drafted an interview script that we read to our users at the beginning of the interview that detailed the narrative role they were playing [Binder Documentation I]; the

first task was highly structured, and involved making travel plans to Los Angeles on a budget of $150 with the following constraints: they live on campus at UCSD, do not have a car, must get to UCLA for their friend's birthday, plan a hike, and take their friend out to dinner at an Indian food restaurant. Our intent with this task was to provide highly focused guidelines so that we could compare the strategies different users employed to reach the same end result. The second task was a slightly less structured shopping task. We simply asked users to browse on the Internet to shop for an item that they planned on buying or wanted to buy. The shopping task also enabled us to compare the different strategies and breakdowns that users encountered when performing the same general task.

The raw data documentation of our second round interviews was similar to our first round; in addition to audio recordings, we made detailed handwritten notes [Binder Documentation I] about how the users approached different parts of the tasks and in which specific ways they managed their tabs. Overall, we noticed that we had each become more skilled in note-taking since the first round of interviews, and found it easier to keep up with the user's actions. Once again, we used the "Master-Apprentice" model and made sure to always ask why the user carried out certain actions or behaviors.

After everyone completed their second round of interviews, we individually analyzed our data on a step-by-step level and visualized it in a table on Google Docs [Binder Documentation I]. In this analysis, we split up the interview into the main categories of "Transportation," "Los Angeles Activities," and "Shopping." For each category, we identified the specific actions that our user performed and the reasons behind their actions. Additionally, we used alternating colors of cell blocks to denote when the user opened a new tab or was

switching between tabs. Although tedious, individually analyzing our interview data streamlined the process of data consolidation and gave us a springboard for writing our affinity notes and sequence models, as we may have overlooked some of the fine details that we were forced to scrutinize over.

After completing our contextual and task-based interviews, we decided to consolidate the data into one comprehensive grid format. First, we made codes for our users based on the first name of the note taker (if it was a paired interview) and the round of the interview. For instance, the first interview of Lara's was coded L1a and the second interview of Natalie's was coded N1b. Although this strategy did not specifically adhere to the guidelines laid out in the book, it allowed us to preserve the anonymity of our subjects and provided us with a succinct way of referring to our interview data. Next, in an effort to consolidated our two rounds of interview data, we condensed our user data into a grid-like format and identified the user's demographic information, the task they completed, task notes, analysis, and any insights we took away [Binder Documentation I]. For the design process, we found it very useful to have a visualization of the main points from our interviews all in one summarized document.

Interpretation Sessions

As in our first round of interviews, interpretations were conducted individually or with two people, depending on team members' schedules.

N2a only did the first task (planned the trip to Los Angeles) because this user started with a very low base level of knowledge regarding transportation, and therefore only was able to complete the first task in the allotted interview time. As for finding a place to hike, the user

searched for hikes near UCLA, and opened four links in new tabs, looked at the content on each of the four pages, and closed each tab before reading the content on the next one. N2a eventually decided she would just let her "friend" decide on a hike to go on. When she wanted to find an Indian restaurant, the user Google searched for Indian food near UCLA, and opened a link for a restaurant guide in a new tab. From the guide page, she opened two links of two restaurants in new tabs. When she decided that she would not want to go to either of these restaurants, she returned to the tab where she had Google searched for Indian food. She then closed the tabs with information on the two restaurants before opening another link for a restaurant guide in a new tab. Again, N2a opened links from the guide page for restaurants in new tabs before deciding to close those tabs (due to "lack" of information on the pages).

D2a exhibited an interesting way of comparing information, as he tended to drag tabs into new windows when comparing information. The user did this for all parts of planning the Los Angeles trip (searching for details on how to get to UCLA, deciding on a hike to go on, and deciding on a restaurant to eat at). D2a also opened links from search results in new tabs and skimmed through each of the pages' content one at a time, but he did so in a very quick, skimming manner.

D2b, as most previous users, opened links from search results/lists in new tabs before making her way through the content of each tab. This user also relied on other users' opinions, as she made some decisions (including where to eat) based on other people's online reviews. For the shopping task, D2a seemed to remember the exact steps to find her dream car, and relied on her memory to re-find what she wanted. A particularly interesting point about this user is that when she opened links in new tabs, she did not close those that were irrelevant or

unhelpful to her. In fact, when her window was “too full” of tabs to continue working through, D2b simply opened a new window and continued her activity from where she had left off in the previous window. A final note about D2b is that she did not seem to know of features (such as keyboard shortcuts) that are “hidden.”

L2a also did not exhibit knowledge of many widely-used keyboard shortcuts. To navigate between search results, she relied heavily on the browser’s “back” button and opening a new tab to begin a new search. She designated tabs for specific sub-tasks (such as finding the name of her ukulele, finding a ukulele case on Amazon, finding a ukulele case on Ebay), and windows for large tasks (the Los Angeles trip and shopping for a ukulele case), but did not close tabs until an entire task was completed (at which point, she just closed the window). L2a also refined search results if she did not see results that she “wanted,” and depended on search autofill and suggestions when she could not remember spelling. Finally, to store “useful” information, she copied/pasted or typed details from pages into her computer’s sticky notes application.

Our next user, Co2a, also did not open links in new tabs, but browsed results by following links within the search window. Similar to L2a, he opened new tabs for separate sub-tasks and did not close tabs until an entire task was completed.

When Cr2a was asked to find a hike for him and his friend to go on, he seemed uncomfortable doing so, and expressed that the reason was that he did not want to use the internet to find a hike, since he did not have prior knowledge of hikes in Los Angeles. Similar to L2a, Cr2a relied on other people’s online reviews to determine which restaurant he wanted to go to in Los Angeles, as long as the content seemed rational and reasonable. As for shopping for a product online, the user navigated directly to his Amazon shopping cart and looked up

product details and checked prices for the product in a Google search before making his decision.

M2a also expressed that he would not use the internet for certain tasks, including finding out how to get to Los Angeles or how to decide on a hike. He claimed that instead, he would normally call a friend who lived in the area and ask for suggestions or opinions. For both finding a hike and a restaurant, M2a opened search results in new tabs before looking at the content on each page and eventually arriving at a decision. When he was deciding on a restaurant in Los Angeles, M2a based his decision on information that was immediately available to him in Yelp! search results (review stars, the number of dollar signs, etc). In addition, the user liked to customize his browser, but not when it required "too much effort."

Our final user, A2a, only completed the shopping task. She used Craigslist and after searching for the product she wanted, she narrowed down the results by price and location. She scrolled through the list and opened links of "potential" items she wanted in new tabs. After she had a handful of new tabs open, the user looked through the content on each of the newly opened pages, and moved tabs that she still wanted onto the left side of the tab dock. closed those that she knew she was not interested in. She repeated the process of opening links in new tabs and closing those she no longer wanted, and eventually ended up with about five pages with information of different products that she liked. A2a compared these last results intensively, clicking between tabs and looking through their content until she made her decision as to which product she wanted.

Affinity Notes

Similar to our first round of interviews, we found it extremely time efficient to write our affinity notes either individually or in pairs during our interpretation sessions. Because this was our second time completing the contextual interview process, the results were much richer. Each team member found it easier to identify and extract relevant data from their interviews and this translated into three times the number of affinity notes that we managed to generate from our first round of interviews.

Unlike the first round of affinity notes, with our second round data we stuck close to the format recommended by our textbook when creating the notes. Each group member studied his or her notes and during interpretation session he or she identified the pieces of data that they wanted to record as affinity notes. We then translated this information into first person statements. Although this may seem like an insignificant deviation from our first round behavior, having a standardized format for the affinity notes made them much easier to compare when we started to build our affinity diagram.

Sequence Models

We found it very useful to create sequence models of our second round data for the main task categories [Binder Documentation II]. Because the users all completed the same narrative task, it was much easier for all of our interviewers to break each task into individual steps. First, each interviewer literally walked through the interview step by step, referencing their notes, and recorded each step in the task workflow into a two-column word document. The left-hand column stored the task information. Next, the interviewer stepped back through each task and color-coded steps that represented breakdowns or triggers. Finally, the

interviewer filled in analysis and breakdowns, triggers, and intents on the right-hand side of the table.

By walking through the interview step by step, the interviewer was able to revisit moments in the interview that might have previously gone unnoticed. In some cases, they identified pieces of information that could be translated into affinity notes and then added to the affinity diagram. In other cases, it was simply an excellent exercise in digesting the volumes of information that was extracted from each interview. In a way, it served as an individual consolidation process that led well into the group consolidation process.

Consolidated Sequence Models

We then broke up into groups and made consolidated sequence models [Binder Documentation II] from the seven individual sequence models for each task. At first we found this process to be quite difficult, and tried a few different formats before deciding on the final one. We made three columns: “activity,” which described the main steps that the users took for each main activity in the interview; “intent,” which described the reason that most users perform this particular sub-activity; and “abstract step,” which detailed the different strategies that users employed to accomplish each specific sub-activity. In this column we also included in red any breakdowns that the users encountered along the way.

Affinity Diagram

Methods

After gathering and consolidating our data from our two rounds of interviews, we began to make our affinity diagram. We took the affinity notes that we created from our first and second rounds of interviews and wrote them on yellow sticky notes. We created this wall in a total of about four different meetings; for the initial sessions, we read aloud each affinity note before sticking it on the wall, and took turns adding to the diagram and placing our sticky notes next to other ones that were similar. We then silently read over the clusters that we had created; as a group, we discussed the creation of each major category as columns formed from the clusters. As a column formed and was finalized, we applied a blue sticky note above it to identify it by category name.

Once all of the columns had been finalized and named, we realized that some of the categories were interrelated and that it would make more sense to add another layer of categorization on top of the current one to show these relationships. Our original blue labels were too concrete, so we stripped them off and replaced them with overarching "I" statements (i.e., in first person). After all the process was complete, we identified seven large overarching "pink label" categories which wrapped our sixteen "blue label" sub-categories. Our seven overarching categories were labeled: "Problems and Frustrations," "Modifications, Customizations, and Desires," "Information Preservation," "Strategies for Accessing Information," "Visibility," "Organization," and "Comparing." We found that going through this process as a group helped us to become more familiar with the data as a whole and to start conceptualizing some of the patterns that were occurring.

Results

The first two categories, “Problems and Frustrations” and “Modifications, Customizations, and Desires”served as our miscellaneous categorizations. The affinity notes classified there tended to be related to specific requests that users made, abstract problems that they acknowledged or voiced, and in some cases , data relevant to a user’s ability to use the browser efficiently. These were the most difficult affinity notes to classify and this was definitely the smallest side of our affinity diagram.

The next two categories, “Information Preservation” and “Strategies for Accessing Information” seemed to encompass the majority of the larger behavioral strategies we identified amongst our users. Due to the fact that our second round interviews were heavily task-focused, we managed to generate many affinity notes that were then classified in these categories. “Information Preservation” covered affinity notes that involved users’ strategies for saving information that they would come back to later. A representative affinity note is, “I like leaving pages open so that I can come back and reference/use them later on.”

The next three categories, “Visibility,” “Organization,” and “Comparing,” were related to the ways that users preferred their content to be represented. “Visibility” contained a large cluster of our affinity notes related to how users prefer their content to be visual. We learned that users prefer to orient themselves visually and liked to be able to visualize system function. “Organization” contained information about how users represent their content spatially. Within the browser, users have adopted a wide variety of strategies to maximize their content visibility and improve their browsing efficiency. In some cases, we found that they used tabs and windows to segment the context of the information that the windows contained. We found

other users that organized the order of their tabs based on different searches or contexts. A representative affinity note from this section is, “I group my information by project.” Finally, “Comparing” contained our affinity notes related to the way users liked to compare content that they accessed within the browser. We were surprised to find that this section was actually rich with examples of different strategies that users employed when they needed to cross reference information. They sometimes resized windows so that they could browse content side by side, sometimes flipped between tabs, and created lists by opening consecutive tabs and then working through them one by one. This section enjoyed a healthy overlap with the previous two sections, as users’ strategies for comparison normally involved spatially orienting their space in a way that made the information easy to visually compare.

Walking the Wall

Once all of our data was consolidated, we began a very thorough “wall walk” process. Adhering to the guidelines detailed in the book, we attached our consolidated sequence models and affinity diagram on the walls and individually read through all of the data from top to bottom. We used smaller sticky notes to identify design ideas and any holes in our data. We kept this a silent process and deferred any questions or comments we had for the group until the end. After everyone finished contributing to the wall, we went through and read each of the stickies aloud, identified which affinity note it referred to, and asked for clarification from the author if necessary. Although time consuming, going through each design idea and hole helped us all to realize which ideas were the most frequently-occurring, giving us a good idea of which ideas to begin pursuing in our design process. After our wall walk meeting, Marie typed

up the entire wall, included every yellow, blue, and pink sticky note and any design ideas or holes that referred to them [Binder Documentation III]. Digitalizing this data proved to be extremely useful for our design process and overcame the issue of our affinity wall and consolidated sequence models only being accessible in one physical location.

Consolidated Takeaways

Upon completion of our second round interviews, our group had amassed and organized an extensive amount of data pertaining to how users from our target demographic initiated, refined, and reiterated searches throughout a task where locating a piece of information was the end-goal. When our team analyzed the completed affinity diagram, it was clear that we had spotted numerous characteristics which we felt provided us with insight into potential design changes that may improve the efficacy with which users employ various search methodology. There were two user traits which we felt provided the most encompassing explanation of why our observed users were doing what we saw, and how we could potentially explain the apparent triviality of some of the breakdowns that had occurred. These two most descriptive observations that our team noted were how our users' Internet search ability led them to preemptively circumvent many problems and fluidly resolve any they encountered, and that our users' apparent model for taking a broad search term and refining it to sufficiently obtain their desired result resembled the structure of a funnel.

First, midway through the design process as we developed design ideas, we found ourselves fearful that the breakdowns we observed were not truly significant enough to

necessitate elements of the process be redesigned. We realized that many of the inhibitions for which we were proposing redesigns did not entirely halt the users' progress, or even appear to meaningfully disrupt it. As a group we discussed this newfound setback and came to the conclusion that our users were experiencing problem-blindness due to their proficiency and natural adaptivity. Because of the experience and expertise of our users, they had become so adept at shaping their process to fit the constraints of the system they had become accustomed to that they infrequently encountered significant disruption. Additionally, because of how skilled users had become at interpreting and rearomsecpmcting to the minor breakdowns they did encounter, they rarely lost a significant amount of productivity when problems were encountered. Ultimately, we decided that while some of the design ideas we had proposed seemed frivolous, they may in fact be much more justified than we could observe from initial interview data analysis alone. Because of our users' ability to adapt to the design elements within the current system (which may have been inhibiting their progress because of its design), our best possible analysis of our design ideas' viability would come from analysis of prototype interview results. Finally, the mere fact that users who possess the skill and experience that our users did were in fact encountering breakdowns at all proved to us that indeed some element of the current system needed redesigning. We concluded that if the current system were perfect "as is," our expert users should have been able to navigate it without encountering breakdowns at all.

Secondly, upon in-depth analysis of our users' search behavior, we found that our initial assumption that users generally had a specific target in mind when beginning a search was not normally true. When we compared all of our user data pertaining to search behavior, we

observed that most users began the search process with an extremely broad variation of what would eventually become their most refined search query, and steadily altered this broad input through utilization of automated assistance features such as autocomplete, through addendums or retractions based on lists parsed using the Curation & Consumption style of results analysis, and through various other refinement methods. We termed this search model the “Funnel” Model because of its resemblance to the common household tool of the same name, which similar to our users’ observed search strategy begins very wide or broad at the start and progresses toward a comparatively small or precise point at the end. We found that this model was exhibited by all of our users during at least one portion of their contextual interview.

Designs

Brainstorming: Our Process

Our brainstorming process took place both individually as well as together during our group meetings. In addition to brainstorming together during our meetings, we took time to silently read through our affinity notes and consolidated sequence models. Then we sat at a table as a group and discussed issues that we had seen. The process was a fluid one—we were constantly challenging our own (and each other's) assumptions as new insights came to light. After we sat down and discussed as a group, our next step in designing would often be to individually flesh out our design ideas. The next day, we would explain our design ideas to the rest of the group, literally grounding the design idea into the data—we held ourselves and each other accountable for this by flagging relevant yellow affinity notes that connected to our design ideas. In this way, we made absolutely certain that our design ideas were derived directly from the data. The affinity notes containing user breakdowns and frustrations were especially useful in our design process.

In the end, we consolidated our data and insights into four main subcategories within the Funnel Model (as outlined in Consolidated Takeaways above). The Funnel Model deals with the way that users behave when they are performing searches on the Internet. The subcategories within it also deal with search behavior, and are called “Curation and Consumption,” “Recognition over Recall,” “Cognitive Offloading and Distributed Memory,” and “Problem Relativity.” All of our design ideas, initial and final, are grounded in these general categories, as well as specific user affinity notes/data**.**

Initial Data-Driven Design Ideas

After our second round of interviews, and upon the completion of our affinity wall and consolidated sequence models, we walked the wall and models, and posted design ideas and holes. After the final completion of our affinity, including these additions, we compiled all of our design ideas and sorted them based on our four main design insights. Our initial design ideas are as follows.

Idea: New ways of organizing that will lessen frustration. Make it aesthetically pleasing and easy to use.

Users exhibited irritation with their own disorganization, as well as with the current organizational system. We wanted to mitigate these frustrations and make improvements. We also noted that our users seemed fond of aesthetically appealing interfaces, especially ones from Google, and decided that making the design of an organizational tool like bookmarks more attractive would make users more inclined to use it.

Affinity notes based on: "I don't organize my tabs;" "I am disorganized and would like to be better at organization;" I like Google products."

Falls under: Cognitive Offloading and Distributed Memory, Problem Relativity

Idea: Make beginner and advanced versions of the interface.

We had a few ideas on how to execute this. The possibility of having an extremely modular, flexible interface that gives the user a high degree of control over the positioning and visibility of elements in their browser would be one solution. We also considered creating a very straightforward basic version of the browser for users who do not wish to customize and have simpler UI needs, and also having a variety of advanced options that allows other users to set up the browser as they wish.

Affinity notes based on: "I use what I know; new tools are hard for me to use;" "I like to co-opt whatever is available to enhance my experience;" (As well as other, almost identical, affinities.)

Falls under: Recognition Over Recall, Cognitive Offloading and Distributed Memory, Problem Relativity

Idea: Give users a "frustrated X" for when they get angry/tired of task but still want to come back later.

We had users who indicated that if they became too annoyed with their task, they would simply give up and close whatever they were doing due to frustration. It was suggested that we add an additional button that allows the user to close a page for the time being, and come back to it later, by preserving it in the same state, somewhere that is easy to access on their computer. This has potential to mitigate switch cost when the user does come back to their task and prevents them from having to start over entirely.

Affinities based on: "If I get frustrated, I will completely abandon my computer;" "I get frustrated and switch tasks if I don't find what I'm looking for quickly."

Falls under: Recognition over Recall, Curation and Consumption (maybe), Cognitive Offloading and Distributed Memory

Idea: A way to prevent accidentally closing tabs. Possible triple click instead of an "X."

Related Idea: Better placement of the "X."

Users displayed some difficulty in finding the "X" to close tabs, and many times accidentally closed the tab when they were trying to open a new tab, since the button to close the rightmost tab and to open a new tab are adjacent. We propose implementing a triple click function instead. Then, there is nothing the user must find, and it is much less likely that a user will triple click something, since this is a slightly unusual action. Alternately, we could simply make the close button more visible, and/or make it modular, so that the user can position it where it easiest for them to use.

Affinities based on: "I don't like accidentally losing information;" "I hate when I accidentally lose tabs;" "I hate accidentally closing things I am using;" "I do not like hidden/poorly marked affordances."

Falls under: Problem Relativity

Idea: Easier way to restore last tab or section button.

In general, users did not have a good way to easily restore browsing sessions that they had accidentally closed or later realized they needed again. Additionally, many users were unaware of features such as Chrome's "Restore Closed Tab." Users also expressed concern in this area about using unreliable computers that had a propensity to crash.

Affinities based on: “I don’t like accidentally losing information;” “I hate when I accidentally lose tabs;” “I hate accidentally closing things I am using;” “I do not like hidden/poorly marked affordances;” “I would like an easier way to restore when my browser/computer crashes.”

Falls under: Problem Relativity, Cognitive Offloading and Distributed Memory

Idea: Make pinning tabs more available.

Very few users were aware of the “Pin Tab” feature that some browsers provide. This feature allows the user to place a tab where they want it, makes a bit smaller, and keeps it open until it is unpinned. We felt that a better implementation of this would be very useful to users, and would facilitate the preservation of information.

Affinities based on: “I hate when I accidentally lose tabs;” “I hate accidentally closing things I am using;” “I do not like hidden/poorly marked affordances;” “I like to leave pages open so that I can come back and reference/use them later.”

Falls under: Recognition Over Recall, Problem Relativity

Idea: Sidebar or something in the corner for utilities like maps, calculators, calendars, and sticky notes.

Our users expressed frustration with not being able to easily access applications and a browser window at the same time, since having both of these open tended to lead to one being obfuscated and/or inactive, both of which caused inconveniences. Users also expressed irritation with the quality of tools such as maps available on various sites. **[See Data Driven Design Ideas]**

Affinities based on: " I would like utilities and applications to better integrate with my browser;" "I would like better access to utilities."

Falls under: Cognitive Offloading and Distributed Memory, Problem Relativity

Idea:A way to temporarily save information, such as bookmarks that expire.

Related Idea: More meaningful bookmarks.

We observed that users who relied heavily on bookmarks for organization lacked an optimal way of organizing them, and also frequently had many outdated bookmarks they no longer needed. We also found that many users kept tabs open for prolonged periods of time to "save" the page. We came up with the idea to create bookmarks which would expire after a period of time (defined by the user), which would create a more secure mechanism than a normal tab to save information, and would address bookmark clutter. Alternately, we discussed creating bookmarks that incorporate a better organizational system so that the user can use them instead of keeping open lots of tabs to save for later.

Affinities based on: "I'm saving information just in case;" "Even if a page of search results did not yield useful information, I will keep it open just in case for possible reference;" "I like using temporary bookmark tabs."

Falls under: Curation & Consumption, Cognitive Offloading and Distributed Memory

Idea: Pop-up branching feature that shows the track you took to get there.

This idea follows the same logical progression as taking the approach of retracing steps to locate lost car keys. We found that many users found it easiest to resume a task or

relocate information utilizing this type of approach. The browser version would incorporate breadcrumbs of the pages in an activity, with possible a mouseover feature to see screenshots of the corresponding pages. It could be likened to a sequence model of the activity history.

Affinities based on: "I like to view content in the order that I open it;" "I use pathway processes to remember sites."

Falls under: Recognition Over Recall

Idea: A highlight tool that allows the user to copy/paste information from a web page to a sticky note specific to the page they are on.

Related Idea: In page highlighting function.

Users exhibited a desire to be able to quickly and easily access relevant information that they wanted from specific pages that they were using. This tool would allow the user to highlight this information while reading through the page, and then when they moused over or clicked on the tab, they would be able to see the information that they had selected for each page. This information could remain associated with the page if the user decided to bookmark the page. Alternately, the information could simply stay highlighted on the page, making it easier to for the user to find, and adding the benefit that it would be easy to locate and preserve multiple pieces of information at once, opposed to tools such as Ctrl + F, which only allow the user to find one thing at a time. Clicking on the information in a sticky note would open up that particular page.

Affinities based on: " I want to be able to save important pieces of information via keeping tabs open or sticky notes;" "Relevant and important information should be easily accessible."

Falls under: Cognitive Offloading and Distributed Memory, Curation and Consumption

Idea: Ability to queue windows that pertain to specific tasks.

Users expressed that they frequently grouped tabs into windows based on a task or project. This would allow for better organization with this behavior, and could function so that the user would be able to see all of the windows open and what they contained with the click of a button , much like Mac Exposé on OSX. This would allow users to keep tasks separated while preserving information.

Affinities based on: "I group information by project;" "I don't like to have too many trains of thought going at once;" I like to leave pages open so I can come back and reference/use them for later."

Falls under: Curation and Consumption, Cognitive Offloading and Distributed Memory

Idea: Design with a Google-like aesthetic.

In general, our users were big fans of Google products, and liked a clean, beautiful aesthetic. We feel that it would be a good idea to maintain such UI standards on any new browser.

Affinities based on: "I like Google products;" "I dislike hidden or poorly marked affordances."

Falls under: Problem Relativity

Idea: Allow customization of button placement.

Users liked the ability to modify their browsers and have control over the presentation of information. We propose that the browser be almost completely modular, so that users have the freedom to put everything where they want it. We feel that this would cut down on confusion, and allow for a great degree of variance in user taste and preference. The first time the user sets up the browser, they would be able to drag and drop everything into place, then lock it in. To modify, they would need to enter an edit mode. This extra step would prevent the user from making accidental changes. Additionally, various presets would be available, and multiple configurations could be saved on the same computer, so that multiple users could all benefit from this feature.

Affinities based on: "I like using everything available to optimize my experience;" "I like to see everything at once to get my bearings;" "I do not like hidden/poorly marked affordances;" "I like to have relevant information easily accessible;" "I place high value on information that is rational and reasonable."

Falls under: Problem Relativity

Idea: Dividers or color coding of tabs.

Users employed a variety of strategies to section their tabs. Some used multiple windows. Some grouped in certain logical patterns. One user created "dividers"- tabs with Google favicons that denoted the beginning of a new section. We thought that this was an interesting behavior, and decided that either the ability to color sections of tabs, or insert dividers for tabs would be a nice addition to the organization of the browsing experience. Additionally, this is a more visual method of organization than simply using titles of pages, and

our users prefered visual over textual information.

Affinities based on: "I organize based on project;" "I like to spatially organize my tabs;" "I like to keep my information segmented;" "I prefer visual over text based cues."

Falls under: Recall Over Recognition, Cognitive Offloading and Distributed Memory

Idea: Keep UI as uncluttered as possible.

Many of our users found clutter to be a huge, distracting nuisance. We feel that minimizing, and if at all possible, eliminating visual clutter from the browsing experience would be extremely useful to our users.

Affinities based on: "Too many tabs open at once distracts me from the task at hand;" "When I am done with a task, I like to close all of my tabs and windows to get rid of the clutter."

Falls under: Problem Relativity

Idea: Keep users aware of system status.

We found that our users liked to know what was going on at all times. They wanted to receive confirmation on whether or not their actions had worked and users also desired clear error messages when something went wrong. Additionally, if something is loading, they would frequently open several tabs, and allow some to load while perusing alternate information. In this case, they preferred some kind of icon on the loading tabs to keep the user aware of the status. We would like to incorporate all of these elements into our ideal browser.

Affinities based on: *From Sequence Models: Liked loading icons, liked concise errors, liked confirmation.

Falls under: Cognitive Offloading and Distributed Memory

Idea: Unique visualizations of information.

Because users preferred visuals, we thought it would be interesting to design a variety of different information visualizations including a mind-map approach andwould convey activity history, that would display parent/child/sibling relationships across pages.

Affinities based on: "I prefer visual over text based cues;" "I use pathway processes to remember sites;" "I know where information is by thinking big to small."

Falls under: Recognition Over Recall, Cognitive Offloading and Distributed Memory

Idea: Based on clicked results, whatever the user is searching with auto-refines.

Users employed the curation and consumption model, and frequently modified the initial query or went through iterations in order to narrow their search to what they were actually looking for. We propose that the user be able to check the boxes of all results that look relevant or interesting to them, and the site or search engine would then refine the results based on what the user marked as important. While this idea is interesting, we doubt the feasibility of this in a browser, but do feel it would be an interesting addition to Google, or sites such as Amazon or Yelp that present results in a similar format.

Affinities based on: "I like to multitask by starting multiple searches so that I can cross-reference information;" "I curate a list multiple times." "I like to be very specific with which search results I click on;" "I want to be able to refine my search results;" "I will open more links from the same parent tab."

Falls under: Curation and Consumption

Idea: Users prefer non-invasive suggestions and self-explanatory systems.

Many users proved to be quite adept at figuring things out on their own, though some did like a little extra guidance. Users did prefer that interfaces be as self-explanatory as possible, and if there are tool-tips or suggestions of any kind, they need to be as unobtrusive as possible, and easy to turn off.

Affinities based on: "I dislike hidden or poorly marked affordances." "I use what I know. New tools are hard for me to learn;" "When I have something I want to accomplish, I will try multiple approaches if I don't get what I want the first time;" "If I can't see how to do something, I won't figure it out."

Falls under: Cognitive Offloading and Distributed Memory

Idea: At the edge of the page, display a small map of all tabs in the window.

Users liked readily available, easily accessible, and visible information. This would allow users to see a small image (facilitating the use of visuals over text) of each tab, and would allow them to get a feel for all of the information that they have available. This is based on a feature in the code editor Sublime Text, which allows the user to see a thumbnail of the entire page along the scrollbar.

Affinities based on: "I like to see lots of information at once;" "I prefer visual over text based cues."

Falls under: Problem Relativity, Cognitive Offloading and Distributed Memory

Idea: Comparison mode. Allows the user to have two "panes" within the window so that they can look at two separate things at once. Similar to the snap to side tool, but within the window.

We observed users frequently needed to either compare details about a specific item to other items on a list, or simply compare multiple pages. We saw users utilize the snap to side tool to do this using multiple windows, and also a user who had two monitors exhibiting the same behavior with the monitors. We feel that having two panes within a given window would give users a more effective way of accomplishing this comparison task, or simply to see more information at a time.

Affinities based on: "I like to see lots of information at once;" "I will often have multiple tabs open containing different ways of searching the same topic;" "I prefer being able to compare things side by side;" "I like to use one monitor for a list page, and one for more detailed pages."

Falls under: Curation & Consumption, Problem Relativity

Final Data-Driven Designs

The last step of our design process consisted of us fleshing out and mocking up the design ideas we believed in the most. From the affinity notes and from walking the data, we had developed strong issues and patterns that were common to our users. The final designs that we developed are all entrenched in these patterns, as well as the overarching "Mental Model of a Funnel," the details of which are outlined in the results section above.

ToolBelt: Easily Accessible In-Browser Utility Bar

Throughout our interview and data analysis process, we noticed that users frequently used a variety of utilities in conjunction with their Internet browsers. These utilities took both the form of both offline applications such as calculators and sticky notes, and online applications such as maps. These users had myriad complaints about the utilities, and even specifically requested better integration of these tools with their browser. Toolbelt is a tool that would help address these user needs. Toolbelt is a sidebar that appears next to the content in the browsing window.

One of the frustrations we found with application use was that when a user is trying to use both a browser and an application simultaneously, something tends to get hidden. When Toolbelt is active, the content in the window shrinks to fit the space it is confined to, so that it is not obscuring anything. Additionally, both ToolBelt and the browser window would remain active at the same time, addressing another user frustration, as users frequently wanted to transfer information between a browser window and an application and could not easily do so since one would be inactive. Additionally, Toolbelt would function in a similar way to the Dock on a Mac—the user would have the ability to set the tool so that it remains open, opens on click, or opens on toggle. This customizability would provide users with the ability to use have ToolBelt only as prominent as they want it, and would prevent it from taking up an undue amount of screen real estate.

Specific insights pertaining to this idea came from User M2a. This user really emphasized how important the integration of applications is to his browsing experience, and

displayed frustration with attempting to use a calculator while also checking investments online. He clicked on the calculator to enter a figure, and it would hide the page he was trying to read information from. He then clicked back to the page to view the information, and it would hid the calculator. Many other users expressed the desire to have some sort of in browser stick note or notepad tool to record information they wanted to preserve for later, as well as a better map tool that is easily accessible, since the maps many sites use are not very effective.

ToolBelt addresses these issues. It provides a customizable (both spatially and in terms of what utilities it contains) interface that allows users easy, simultaneous access to both their browser and desirable utilities. It removes the issue of hidden information, and allows both things the user is using to be active at the same time. It additionally facilitates cognitive offloading and distribution of memory, as the user is now able to quickly and easily transfer and preserve information utilizing this external entity.

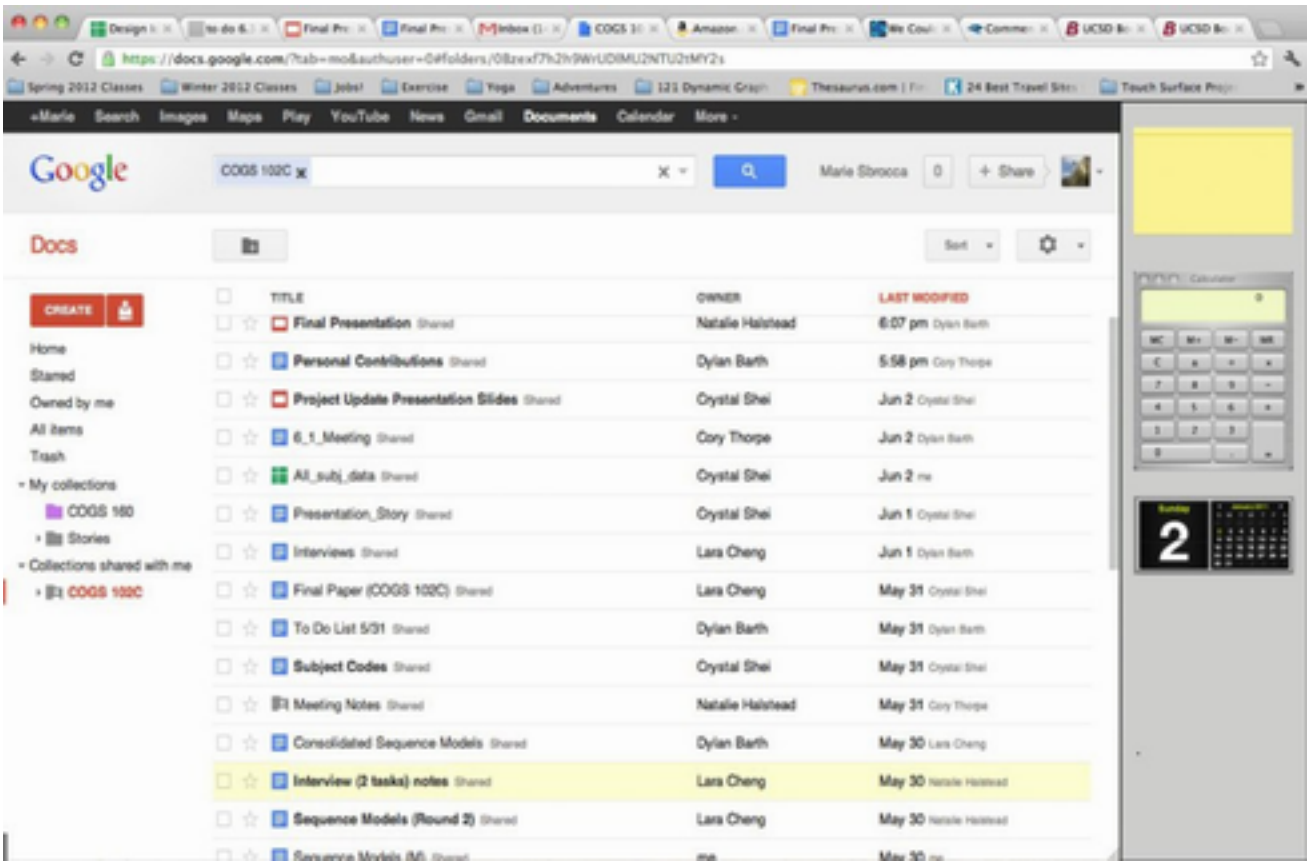

Figure 1. Toolbelt Mockup. Toolbelt appears to the right of the content displayed in the browser, while it is open. It contains a set of utilities specified by the user, and its display settings can also be modified to fit the user's preferences.

Affinities based on: "I would like applications and utilities to better integrate with my browser." "I would like better access to utilities;" "I like to see lots of information at once;" "I like to see distance, not just text, on a map;" "I need to have relevant/important information easily accessible;" "Before closing tabs, I write down or record necessary information."

Falls under: Cognitive Offloading and Distributed Memory, Problem Relativity

LightScroll: A Visual & Hierarchical Comparison

We garnered the visual aesthetic of this tool from the scrolling filmstrip preview of a commonly used photo processing software from Adobe. LightScroll, as we call it, is a tool that helps users categorize results when they are searching for things. It most useful when users are in the "Curation and Consumption" mode of their search. The user opens up what we will refer to as the "base"—a page with a sizeable list of links (e.g., Google search results).When a link is "opened as a base in LightScroll,"(refer to *Figure 2.)* a mini screenshot/title of that link's page will appear in a fixed list-type scrollbar at the bottom of the base list. Users can click on a mini screenshot/title in the list-type scrollbar. This will open the tab as a page with a tab next to the base page. The list-type scroll bar will be fixed on every page that is opened from one of its tabs (e.g. all tabs within the same base grouping will have the same "LightScroll" display). Pages can be closed by clicking on the X button on EITHER the screenshot in the list-type scrollbar, or simply by closing the tab (as is currently conventional). In this design, the children tabs remain grouped on the bottom of every page that has the same "base" group, and their respective pages remain unopened until their tabs are clicked on. The base tabs will appear in the traditional tab area, at the top of the page, while their children tabs will only be visible when the base tab is active. In this way, we encourage the users to self-organize their trains of thoughts.

The insights that led to this design came from many users' data. For example, as seen in L1a's interview, the user's reasoning for looking through the group of newly opened tabs after making his way down a starting list (rather than making his way down multiple pages worth of lists) is that with too many tabs open, the tabs become hard to click and harder to read, as the

titles become visually occluded. L1a was searching through videos, so he had to manually pause each video in each new tab. This system would address that issue as well, as the LightScroll would not load the page until the user indicated through some action, either by clicking the tab or manually starting it.

In addition to addressing clutter reduction and visual organization, the nature of the LightScroll tool is such that users with different preferences could all potentially still benefit from its use. Users who like starting many trains of thought at once (see affinity note below) can start as many trains of thought as they want (in the form of the base tabs). The parent/child model of LightScroll will help these multiple trains of thought users to structure their thoughts more clearly. The users who do NOT like starting many trains of thought at once could be less intimidated by multitasking; if there is built-in structure and activity context, these users could branch out from their comfort zone. And, if they still choose not to use multiple base tabs at once, they still benefit from the visual aid that the LightScroll tool provides.

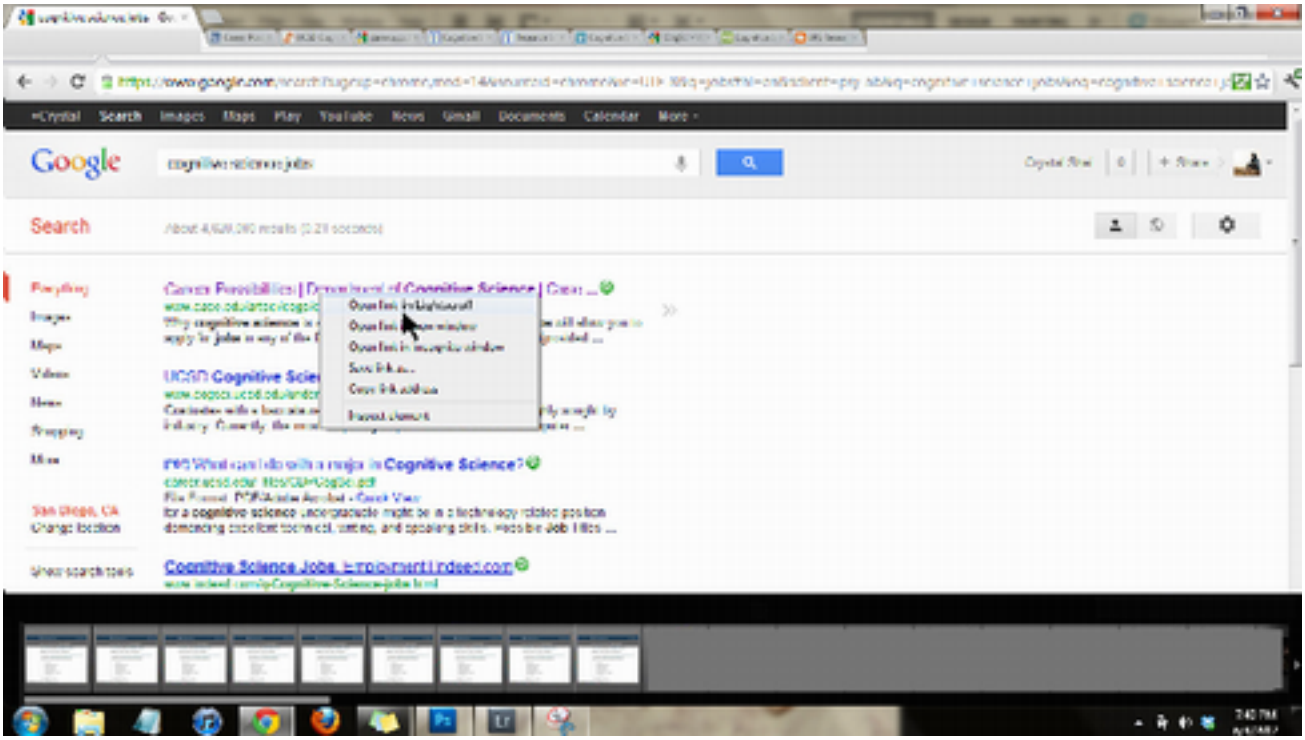

Figure 2. LightScroll Mockup. LightScroll appears on the bottom of the screen. LightScroll can be opened with a right click option or a keyboard shortcut. The "base" tab, or parent tab, appears on the top left corner. The children tabs appear smaller to the lower right.

Affinities based on:"I don't like clutter;" "I like to organize spatially;" "I don't like having too many tabs at once;" "I don't like starting many trains of thought at once;" "I like starting multiple trains of thought."

Falls under: Curation and Consumption

FootPath: Breadcrumb Visual History that allows a whole "path" to be bookmarked

The FootPath idea is a fusion of regular tab bookmark abilities and website breadcrumbs. The design would consist of a bar below the URL bar in the browser, with a real-

time visual representation of the user's chosen path. As the user browses through a path, the FootPath would update with titles and screenshots, though the screenshots would only appear on mouseover (see *Figure 3* below). There would also be toggleable options like "Bookmark This Path" or "Start a New Path" to allow the user flexibility in using the tool. If the user does not wish to turn it on except in specific cases (e.g., if they are working on a research project) they will have the ability to do so. The idea behind FootPath is that it would be a useful tool during a search for a project or a work period. Perhaps a user will come across a piece of information that they only realize is relevant hours or days later. With the FootPath enabled, that piece of information will not be lost. In fact, FootPath will allow the user to easily follow his or her thought process when searching. An extension of FootPath that we discussed but did not implement into our mockup is: a replayable record of browsing history in which a user can select a day or time range, and then watch the browser history unfold in real-time, or various degrees of fast-forward, accompanied by small bird's eye view of the pages visited. Incorporating activity history into the browser history would provide users with the visual context they have expressed a desire for, and the history visualization they are accustomed to from the path bookmarking.

FootPath is grounded in the findings of our user populations. Two specific users, Cr2a and D2a, both used the "pathway" method of re-finding a site that they had forgotten the direct URL to. In both their cases, they had previously accessed a site by jumping through various other sites. In order to come back to that end-point site, they jumped through the same steps, much like a virtual retracing of steps. Other findings that FootPath is grounded in include the users that utilize "temporary bookmarks," as this would eliminate the need for those, as

well as cater to those users that like to organize their bookmarks/tabs spatially.

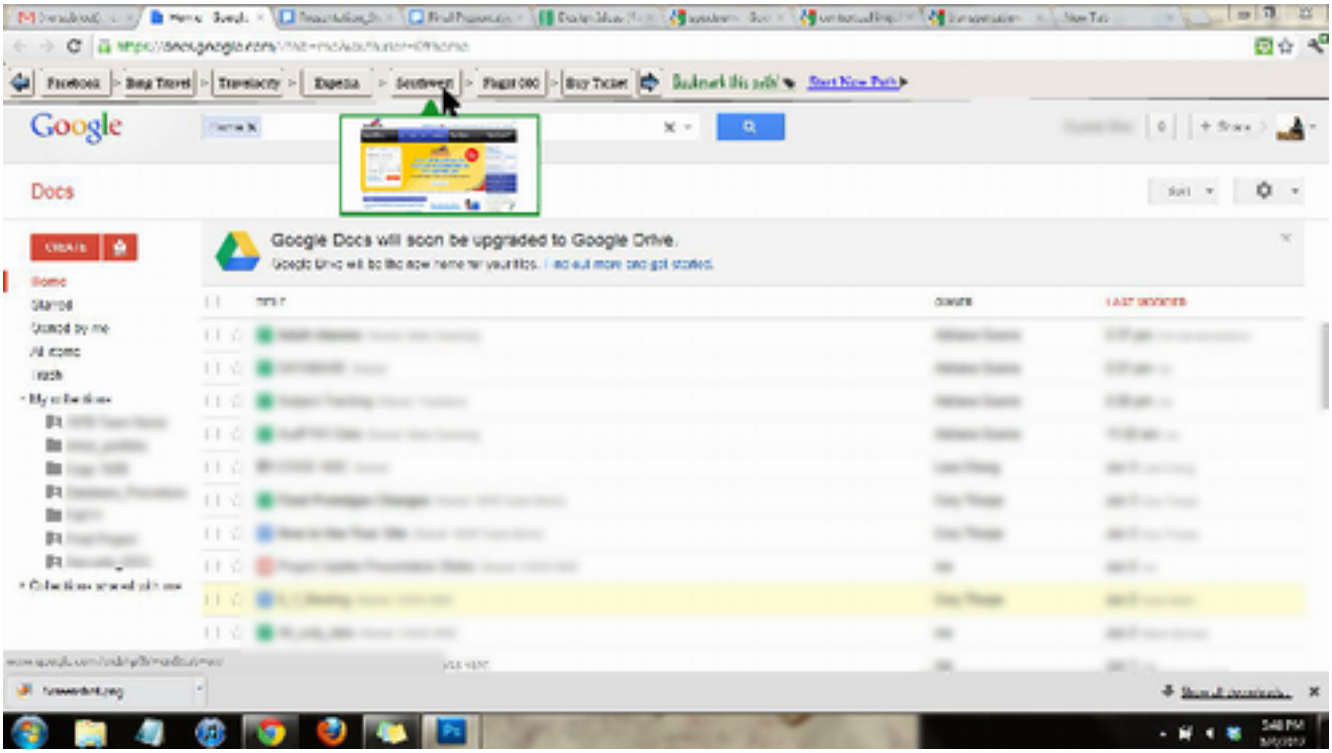

Figure 3. FootPath Mockup. Cursor hovers over one point in the path and a screenshot of the site appears. A button arrow option to navigate the path is on the left/right sides of the path. User also has the option to bookmark the path or start a new path.

Affinities based on:"I use pathways/processes to remember websites I need;" "Sometimes I need to come back to a site but did not bookmark it;" "I have too many bookmarks;" "I hate accidentally losing information;" "I like to organize spatially;" "I like to leave pages open-I'm saving information just in case;" "I will make temporary bookmarks, and then come back later and delete"

Falls under: Recognition over Recall, Cognitive Offloading and Distributed Memory

Discussion

This project was a multifaceted endeavor with many frustrations and setbacks. Balancing the project with our pre-existing obligations from our other classes and work all within the constraint of the ten week quarter system was certainly difficult for everyone involved. However, our project did ultimately prove to be extremely rewarding and we came out of it with a greater understanding of both the way Internet power-users structure information in their minds while browsing and of the general process and power of data-driven design.

As far as studying college students and their usage of tabbed browsing goes, we made dozens of interesting observations and developed several different insights about that particular behavior. One of the most significant of these insights was our "funnel" mental model. This model states that users access and search for information in a manner reminiscent of a funnel: An initial broad sweep representing the top of a funnel is narrowed down continuously until the user is left with a very small yet very pertinent amount of information. In this model, users interact with the facilities available to them in the system to search, refine, and compare results to eliminate weaker choices and select stronger ones.

By using this mental model as a jumping off point, we were then able to further identify specific behaviors among our tested users: Curation and Consumption, Recognition over Recall, Cognitive Offloading/Distributed Memory & Cognition, and Flexibility & Problem Relativity. With these behaviors in mind, we then drafted design ideas with the purpose of using them to benefit the users in performing these behaviors. We took the best of these design ideas and

extensively refined them to come up with our final list of design ideas: ToolBelt, FootPath, and LightScroll.

On a broader scale, we learned many things about the design process and all of the smaller subsets of activities that make up that process such as conducting user interviews, interpreting the interview data and creating design ideas. One of the most significant lessons we learned was the importance of staying grounded in the data. When coming up with our design ideas, we were very eager to come up with interesting features as well as ways to completely change the modern browser. While we found it interesting to come up with these wild ideas, at the end of the discussion the question was always asked: “What problem is this actually solving?” Frequently, the discussion of an exciting new feature would actually have no real purpose in solving a user breakdown or streamlining a workflow. Situations like these served as a constant reminder of the importance of staying grounded in all of the data that we collected.

Future Plans

The short timeline imposed by the quarter system meant that we had to not only set up a strict timetable for our group but also to abide by that schedule very closely. While we did establish a detailed timeline early on in the quarter, we ran into a significant hurdle when we realized we needed to conduct more user interviews to draw additional data from. We had to immediately prioritize doing that and thus were forced to pushed back the entire design process from that point on.

With more time to continue the project, we would continue the contextual design process by creating paper prototypes of our design ideas and then conducting paper prototype interviews to refine our designs. In addition to this, we would have also liked to gather data from a broader range of users. We targeted power-users from the start of our project. However, gathering data from a more diverse range of internet users with differing backgrounds, jobs, skill levels, and ages would have been an excellent opportunity to use the insights we gathered from our pre-existing interview and to extrapolate from these insights to the wider range of users. Finally, we would have liked to talked to more people with experience in design about our project. Speaking to people like Professor Hollan, the TAs, and the professional UX engineer was by far one of the most helpful parts of the project. Given the chance, we would like to consult such professionals to show them our data and receive feedback on how we would implement our findings and design ideas.