Understanding How Users Map Regions Between Web Pages

Salman Ahmad
Stanford University HCI Group
Computer Science Department
Stanford, CA 94305
saahmad@stanford.edu

Ranjitha Kumar
Stanford University HCI Group
Computer Science Department
Stanford, CA 94305
ranju@stanford.edu

ABSTRACT
To assist novice users in creating web pages, we envision a tool that automatically transforms an existing web page into the layout and design of another. In order to achieve this vision, one must develop a learning algorithm that is capable of mapping regions between web pages. To do so, an in depth understanding of user behavior is necessary. Thus, we seek to understand how users map regions between web pages. There were two fundamental questions that were asked: (1) do users map web pages consistently and (2) what motivates their mapping decisions? A custom interface, which asked users to select corresponding regions between two web pages, was utilized in a Mechanical Turk study and in a lab study. We found evidence that indicates that users do map web pages consistently and discovered that users tend to create mappings in a manner that preserves the hierarchy of the pages while also pairing semantically salient elements.

Keywords: web page design, crowd source, user study

INTRODUCTION
For most users creating web pages is difficult. There are numerous technical hurdles and design challenges to overcome. The majority of tools available are aimed at abstracting technical details from users (for example, WYSIWYG HTML editors) but most offer very limited affordances for design exploration (usually in the form of limited, predefined templates). This is truly unfortunate because the number of web pages on the Internet has opened up new horizons for leveraging examples. Examples can very helpful by allowing users to explore existing solutions and also get inspired [1]. We envision a tool that enables a user to seamlessly use any existing web page as a design template, allowing them to automatically transform the content of an existing page into the layout and style of a target page.

The pivotal component in such a system is the mapping algorithm to automatically maps regions between web pages in a manner that mimics human behavior. This of course makes the assumption that a good mapping is equivalent to a human mapping. Since web pages are designed to be viewed by human users, this assumption seems reasonable. Before such an algorithm can be created it is important to understand how human users map regions between web pages. With this context, we sought the answer to two questions:

Do users consistently map web pages? Knowing this allows us to validate the fact that developing a learning algorithm is in fact possible and worthwhile.

What are the factors that cause users to make mapping decisions? Knowing this will allow us to build intuition on how to build and optimize the algorithm to deliver good results.

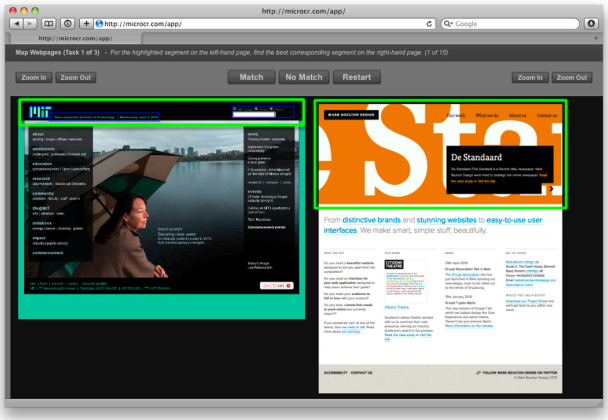

Figure 1: An overview of a system that allowed users to automatically leverage existing examples by transforming a source page into a target page.

To explore these questions a custom interface was created to ask users to map regions between web pages. This interface was used in conjunction with two studies, one using Mechanical Turk [2] as a crowd sourcing platform and the other in an in-lab study where users were observed and asked to think out loud. In this paper we describe the details of our studies and summarize the key findings.

RELATED WORK
Due to the highly specific nature of this study, most related work focuses on completely different questions. Most similar studies fall generally into two categories, analyzing single pages for accessibility reasons, marketing purposes or spam detection. In terms of creating a tool that leverages existing examples, Fitzgerald created a tool that allows users to transfer CSS styles between pages. This is, however, limited to copying specific style rules [3].

USER STUDIES
As aforementioned, this study aims to explore two questions, do users map content consistently and what are the factors that motivate their mapping decisions. To answer these questions two user studies were designed. Users were shown a graphical interface for selecting mappings (described below) and simply asked to follow their intuitions when selecting corresponding regions.

The Mechanical Turk study, due to the large sample sizes that were possible, was designed to primarily answer the first question. Overall, we received 59 good HITs and paid users $0.50 each. The lab study was performed to get a better idea of factors that motivated users to map regions. Three users were recruited for the study and asked to think out loud when selecting mappings. Both the audio and screen activity of these sessions were recorded and after the task was completed, users were given the opportunity to reflect on their performance.

In both studies, the same user interface was used. This interface presented users with two pages at a time (Figure 1). The pages were decomposed into visually salient regions using a highly modified version of the VIPS [4] algorithm. A particular region on the left side page was first chosen and users were asked to select what they felt was the corresponding region on the right side page. A user could alternatively select "no match" if they felt no match existed. Periodically, users were asked to explain why they created a particular mapping pair by means of popup dialog window. This was particularly crucial for the Mechanical Turk studies where we could not observe user behavior.

All users were asked to map the same three pairs of pages (six unique pages in total). These pairs included a pair that was structurally similar, structurally dissimilar, and a pair of pages that were familiar to users.

RESULTS AND DISCUSSION

There are threes terms that are frequently used in our discussion, "pair", "mapping", and "similarity". To avoid any confusion, we define a "pair" as two regions (one from each page) that a user felt corresponded with each other. A "mapping" is a collection of pairs created by a user that summarizes which regions from one page map to the other. "Similarity" between two mappings is the number of shared pairs between them divided by the total number of unique pairs. We also remind users that there were three pairs used, one similar, one dissimilar, and one that was familiar to users.

The results show a global average pair-wise similarity of 61.9%. A full break down of the results can be seen in Table 1. Looking at the results, what is immediately obvious is that users agreed more on the pages that were structurally similar. However, some of the numbers seem suspicious. While it is promising the in some case 75.2% similarity was achieved, numbers like 46% are seem hardly convincing. It is important that readers remember that the aim is to see if there is consistency in the result such that a learning algorithm can recognizes patterns from human data. The fact that users agreed 46% of the time in the worst case is promising. Furthermore, it is important to note that a huge number of the mappings were off by just one or two pairs, a fact that is hidden by the raw percentages.

	Similar	Dissimilar	Familiar
Mean	46.0%	75.2%	64.4%
Median	47.1%	85.7%	64.0%
Mode	50.0%	92.3%	66.7%
Std. Dev.	18.4%	23.3%	14.9%

Table 1: Summary of results. Values are the pair-wise similarity percentages.

Another interesting take away that is not shown in the summary statistics, is that in each of the three mapping tasks, there existed pairs that every user agreed on. These tended to fall into two categories: hierarchical elements like headers and footer and then elements with strong semantic meaning like logos, navigation links and search boxes. Analysis of the user submitted reasons and observations from the lab study confirm these findings. In the labs study, we observed that all users attempted to preserve the hierarchy of the site. However, while doing so, if they came across elements that had strong semantic meaning, they would readily break the hierarchy and map to what they felt was the semantically corresponding block on the other side. This was particularly evident with a mapping between Bing.com and Google.com, where users would all preserve hierarchy until they were asked to map the logos and the navigation links, which appear in different places of the hierarchy in both pages.

Overall, our study did find evidence that users map web page in a consistent and similar manner. Furthermore, our study shows that users tend to be more consistent with hierarchical mapping (since the structurally similar page had a much higher similarity percentage) although are more likely to create a mapping pair between semantically related regions (since all users consistently agreed on mapping pairs based on semantic meaning even though they broke the hierarchical structure). As a result, when building the mapping algorithm, it appears best to focus on attempting to preserving hierarchy while not forgetting to take into account regions with semantic meaning (such as logos and form elements).

REFERENCES

1. B. Buxton. *Sketching User Experiences: Getting the Design Right and the Right Design*. Morgan Kaufmann, 2007.
2. Amazon Mechanical Turk. http://www.mturk.com
3. M. Fitzgerald. Copystyler: Web design by example. Technical report, Massachusetts Institute of Technology, May 2008.
4. D. Cai, S. Yu, J.-R. Wen, and W.-Y. Ma. Vips: a vision-based page segmentation algorithm. Technical Report MSR-TR-2003-79, Microsoft Technical Report, 2003.