

Using card-sorting techniques to organize your intranet

Whether you're designing an entire intranet or just one site within it, one of the biggest challenges you'll face is how to organize the information you want to make available. Organization that is intuitive and logical to your users is crucial. If the users can't find the information they need, it might as well not be there.

At Fidelity Investments, our intranet has changed a great deal over the eight years or so that it has existed. It has evolved from something that only the "techies" used or cared about to something that virtually all employees use every day to get their jobs done and ultimately provide better service to our customers. It is critical to our business and thus critical that our employees can use it effectively.

Back in 1997, we started a corporate-level effort to redesign the intranet; at that time, it was a loosely connected set of internal sites with very little organization. We wanted to organize the material in a way that would make sense to our employees. As the "human factors" or usability member of the team tasked with this redesign, I was asked how we could learn what organization would make the most sense to the users. That was when I remembered a technique I had used many years earlier, in 1985, when working at Burroughs Corporation, a mainframe computer manufacturer. The challenge back then was to determine how to organize hundreds of mainframe operating system commands into a meaningful menu hierarchy. We used a card-sorting technique then.¹ Even though the technology had changed dramatically over the intervening years, it still sounded like a good idea.

After this first major redesign of Fidelity's intranet, we used card-sorting techniques many more times – for the entire intranet, for many individual intranet sites, and for our external sites. As the prototypical example, I'll use that

first redesign to describe the basic technique. Then I'll cover some of the variations we've used over the years, including online card-sorting techniques.

Setting up a card-sorting exercise

The first step in any card-sorting study is to develop a list of the topics that will become your cards. Typically, these topics represent the key pieces of information that your site will contain. They usually correspond to possible content pages, excluding any pages that are purely for navigation. If you're designing an entirely new site, focus groups of users can help develop a list of topics they would like to see. In our case, we were primarily redesigning an existing set of sites, so our list was basically an inventory of the key topics, plus some new ones we knew we wanted to add. After combining topics that were extremely similar, we came up with a list of 90 topics, including the following examples:

- Expense report policies
- Fidelity annual reports
- Intranet usage policy
- Anti-virus information.

In addition to the short titles for each topic, we developed a brief explanation of the topic – no more than one paragraph – and assigned each topic a unique number. We found it helpful to record all of this in a spreadsheet.

From a practical standpoint, we've found that it's difficult to do card-sorting studies with more than about 100 topics. The reason is simply how much time most users are willing to devote to the task of sorting the cards. In our experience, about one hour is the most you can usually expect someone to spend. Although it varies greatly depending upon the user and the complexity of the topics, somewhere around 100 cards is about the most someone can handle in an hour. This means that you might have to redefine the scope of your effort, or the level of detail of your cards. One technique I've used for reducing the number of cards is to ask a few particularly "willing" users to go through your large stack of cards and sort them into closely related groups. Any cards that all of the users put into the same group can be combined into one and described appropriately.

Once your final list of topics and descriptions is set, print them on standard index cards. Include the identifying number – we often put that on the back of the card. You can either use a low-tech approach, with hand-written cards, or the high-tech approach that involves a mail-merge from a spreadsheet into a word-processing document set up to print on perforated sheets of index-card stock. Sometimes it's helpful to make multiple sets of index cards, which makes the high-tech approach more efficient.

How card sorting works

At this point, you need to recruit prospective users to do the card-sorting task. They need to be representative of the target users of your site. The smallest number of participants we have ever used is about 10. Whenever possible, we try to get at least 15 to 20. In the card-sorting study for our initial intranet redesign we had 26 Fidelity employees do the sorting. Logistically, if you have multiple sets of cards, you can have more than one person doing the card sorting at a time. Sometimes we go to the users and sometimes we ask them to come to us.

In addition to the individual cards with the topics on them, participants are also given small white envelopes, larger brown (manila) envelopes and two types of blank index cards, for example, two different colors. (See figure 1.) Typically, the card sorting is done at a desk or table – just make sure there's enough room to spread out the cards! The cards are shuffled so that each user receives them in a random order. Users are asked to sort the cards into logical groupings using no more than a two-level hierarchy. Their instructions are as follows:

- Sort the related index cards into groups. Cards do not have to be grouped if you do not see any relationship to other cards. If you believe an index card belongs in more than one group, use a blank index card to duplicate the card. Use the blue cards to add any new pages/sites that you would like to see.

- Place each "group" of cards into a white envelope. Label each white envelope with a name that describes the information it contains.

- Place related groups of white envelopes and any related individual cards into the larger brown envelopes to create "categories." Label each brown envelope with a name that describes the information it contains.

- Place everything into the folder labeled "Home Page." This folder can contain individual cards, white envelopes, and brown envelopes.



figure 1

Each participant in a card-sorting study is given index cards containing descriptions of the topic areas for the site and asked to sort them into logically related groups, which are put into the white envelopes. Those groups are then put into logically related categories and placed in the larger brown envelopes.

Note that we used a two-level sort for this particular study. With fewer cards (e.g., less than 40 or so), it may only be necessary to use a one-level sort.

Analyzing the results

For each participant in the study, a matrix, or table, is created that reflects that person's perceived "distances" among all pairs of cards. (See figure 2.) The distances are defined as follows:

- If two cards are in the same group, the distance = 0.
- If two cards are in different groups but in the same category, the distance = 1.
- If two cards are in different categories, the distance = 2.
- For duplicate cards, the lowest distance is used.
- Anything at the home page level is treated as a "category" for analysis purposes.

These matrices are then averaged across all participants, and the matrix of average distances submitted to two statistical analyses:

| Card # | 1 | 2 | 3 | 4 | 5 |
|--------|---|---|---|---|---|
| 1 | 0 | 1 | 2 | 2 | 1 |
| 2 | 1 | 0 | 1 | 2 | 1 |
| 3 | 2 | 1 | 0 | 1 | 1 |
| 4 | 2 | 2 | 1 | 0 | 0 |
| 5 | 1 | 1 | 1 | 0 | 0 |

figure 2

Sample of a portion of a "Distance Matrix" for one participant. The diagonal line through the table indicates that the distance between any card and itself is 0. In this example, the participant put cards "4" and "5" in the same group (distance = 0), cards "1" and "5" in different groups but the same category (distance = 1), and cards "1" and "4" in different categories (distance = 2).

- **Multi-dimensional scaling (MDS) analysis** creates a two-dimensional map that depicts the average distances between the cards as closely as possible. (See figure 3.)
- **Hierarchical cluster analysis** merges the objects into clusters based on their distances. The clustering continues until all of the objects merge into a single undifferentiated group. (See figure 4.)

These analyses can be done using several commercially available statistics packages, including SPSS and NCSS.ⁱⁱ

In our initial intranet study, the average number of objects (cards or envelopes) placed on the home page was 10 (range: three to 21); therefore, we started analyzing the clusters by taking a "10-cluster slice" through the data. This means that we looked for a place in the hierarchical cluster chart where we could draw a vertical line that would cut through 10 lines (groups).

In interpreting these results, we find it helpful to compare and contrast the MDS results and the hierarchical cluster results. The idea is to try to detect the groupings of cards that they reflect. See if you can interpret those groupings and assign logical names to them. Study the names that the users gave to their groupings and try to use the same terminology whenever possible. These apparent groupings can then be used to design the top-level representation of these groups on your home page as well as the next level. A great deal of usability evidence has shown that you should generally strive for as few levels of pages in your site as possible – i.e., strive for "breadth" over "depth."

Card sorting for groups

An alternative to the traditional card-sorting technique is "affinity grouping," which uses groups of users to perform the sorting. In this approach, instead of printing the topics on index cards you would print them on "sticky notes"

that can be put on a whiteboard or series of flipcharts. As with the index cards, you would initially array the topics in a random order. A group of users works together to rearrange the topics into the groupings that make sense to them. The instructions are similar to those used in the individual card sorting. Some method of naming the groups (e.g., using a different color of sticky notes, or writing the name on the flip-chart) and identifying any second-level grouping used is given.

The group approach works best with relatively small groups, about six to eight users. Larger groups tend to have difficulty reaching consensus. We generally run at least three or four independent grouping exercises. Although it's possible to run the same statistical analyses as you can with the individual card-sorting approach (treating each group the same as an individual in the analyses described before), we generally just "eye-ball" the results from each group, trying to detect similarities and reconcile any differences. One advantage of the group approach is that you can usually get the results quicker. A possible disadvantage is that small groups are sometimes easily influenced by one or two particularly persuasive individuals, thus possibly yielding results that are not truly representative of the larger target audience.





figure 3

A multi-dimensional scaling (MDS) chart of the average distances between the cards. You can start to see groups of related cards. Each circle represents a card. For example, the tight cluster of circles on the far right is all of the cards representing Fidelity business-unit home pages.

If your intranet or site has different types of users that you think might have different perspectives on how the information should be organized (e.g., technical vs. non-technical personnel), it can be helpful to do separate group exercises with each constituency. Then you can compare the results to see if they really do have different views of how the information should be organized. If so, you should consider offering them different views of the site (i.e., different home pages). Of course, this can also be done using the individual card-sorting approach by simply

analyzing the data from the different types of users separately.



Online tools for card sorting

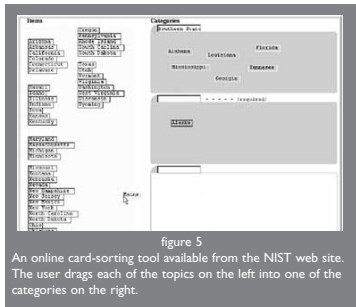
Recently, some online approaches to traditional card-sorting techniques have become available. Two such tools are EZSort, which is available from the IBM Ease of Use web site,ⁱⁱⁱ and WebCAT, which is available from the NIST web site.^{iv} EZSort is a Windows-based application that the user must install. WebCAT is a web-based application that has to be installed on your web server, but which any user with a web browser and access to your server can then use. Otherwise, these tools are relatively similar.

Figure 5 shows the interface for "online card sorting" that is presented to the user by WebCAT. The user can "drag and drop" each of the "cards" on the left into a group on the right. Groups can be pre-labeled or you can require the user to label them. There are certain limitations on the number of topics and how long the description of each can be. It currently only supports one level of grouping. After the user has sorted all of the topics into groups and submitted, the data is automatically tabulated. At any time, the designer of the study has access to an online analysis tool that does a hierarchical cluster analysis on the average of all the data submitted up to that point. It even provides a mechanism for interactively visualizing different "slices" through the clustering. MDS analysis is not currently provided.

A great advantage of the online approaches is that you can potentially get data from large numbers of users very quickly. In a recent online study we did as a part of redesigning my department's intranet site, we got 381 people to do the online card sorting in just two days! Some people argue, however, that there's something to

be said for the low-tech approach with physical index cards. We've never tried to do a direct comparison of the two techniques to see if we get different results.

Regardless of the exact technique you use, card sorting is an excellent technique for learning how your prospective users view the relationships among the various topics of



your entire intranet or site. The resulting design should be one that allows your users to find the information they're looking for quickly and easily. ◀

ⁱ Tullis, T.S. (1985) Designing a Menu-based Interface to an Operating System. *Proceedings of CHI'85 Conference on Human Factors in Computing Systems*, San Francisco, CA, April 1985.

ⁱⁱ SPSS: <http://www.spss.com/>

ⁱⁱⁱ NCSS: <http://www.ncss.com/>

^{iv} http://www-3.ibm.com/ibm/easy/eou_ext_nsf/publish/410

<http://zinc.nist.nsl.nist.gov/WebTools/WebCAT/overview.html>

Thomas S. Tullis is Senior Vice President of Human Interface Design at Fidelity Investments, where he manages a department within Fidelity's Center for Applied Technology. He has more than 27 years of experience in the human-computer interface field. In addition, he has published more than 35 papers in technical journals and has been an invited speaker at national and international conferences. He holds eight U.S. patents.

Prior to joining Fidelity, Tullis worked at Canon Information Systems, McDonnell Douglas, Unisys Corporation and Bell Laboratories. While at McDonnell Douglas, he helped design the human-computer interface on board the International Space Station now being built. Tullis is active in the Human Factors and Ergonomics Society (HFES), the Special Interest Group on Computer-Human Interaction of the Association for Computing Machinery (ACM SIGCHI), and the Usability Professionals Association (UPA). In addition, he serves as a reviewer for various journals in the field, and is chair of the Computer System Technical Group of the HFES.