

EVALUATING INFORMATION ARCHITECTURE

A PRACTICAL GUIDE TO
ASSESSING WEB SITE ORGANIZATION

BY STEVE TOUB, ARGUS ASSOCIATES

NOVEMBER 2000

CONTENTS

INTRODUCTION.....	2
EVALUATING INFORMATION ARCHITECTURE: AN OVERVIEW	7
EVALUATING STRUCTURE	14
EVALUATING GROUPING.....	18
EVALUATING LABELING	23
ENDNOTES.....	24
ABOUT THE AUTHOR.....	26
ABOUT THE ARGUS CENTER FOR INFORMATION ARCHITECTURE.....	27

INTRODUCTION

Information architecture (IA) is the art and science of structuring and organizing information environments to help people effectively fulfill their information needs. This white paper explores the why's, what's, and how's of evaluating a web site's information architecture. It aims to raise consciousness about the evaluation of IA and to provide:

- Web site owners and other decision-makers with an understanding of evaluation issues; and
- Information architects with a synthesis of evaluation techniques.

This paper is not the last word on the subject. Rather, it aims to assemble existing knowledge on evaluating information architectures in one place as a springboard for further advancing this evolving field.

It will:

- Provide background on why evaluating IA is important.
- Place IA evaluation in context and provide an overview of methodological issues.
- Describe techniques for evaluating aspects of an IA, particularly structure and grouping.

THE IMPORTANCE OF INFORMATION ARCHITECTURE

Flaws in information architecture can affect a web site's bottom-line quite dramatically. This impact is most applicable to **e-commerce sites** (also known as B2C or Business-to-Consumer sites).

- **Lower revenue caused by failed searches:** One consulting firm observed that 56% of search engine queries on e-commerce sites ended in failure and claimed that this translates into a loss of billions of dollars.¹
- **Lost revenue caused by abandonment and attrition:** One market analyst report states that when looking for products, 62% of online shoppers had given up at least once. 42% had turned to non-Web channels to complete their purchase.²

While the impact of flawed information architectures on **intranets** may be harder to measure in monetary terms, the effects are no less significant:

- **Lost productivity caused by not knowing where to look:** At one IT company, 41% of employees surveyed stated that they were unsure of what channel of information (e.g., email, intranet) to use to locate knowledge.³
- **Lost productivity caused by not being able to find needed information effectively:** One oil company's head of information services reported that 35% of an employee's time is wasted looking for information and that over the last 10 years that number has been as high as 50%.⁴
- **Lost productivity caused by not being able to manage content effectively.**
- All of these inefficiencies can lead to decreased employee morale, compounded productivity losses, and can contribute to unnecessary turnover and training costs.

The impact of flawed information architecture may also have a significant **impact on information architects themselves.**

- **Legal liability.** Consider IAM.com's recent lawsuit against Web development firm Razorfish for delivering a site that was "flawed by grave technical and navigational problems."⁵ We now live in a world where information architects can be sued for poor IA design.

Once decision-makers understand the importance of IA, the next logical step is to want to measure the results of IA efforts.

CURRENT PROBLEMS WITH IA EVALUATION

Given the importance of a well-designed architecture, it is surprising how few companies evaluate their web site architectures. For example, in a survey of executives who are responsible for the searching function of e-commerce sites, analysts found that although 90% of respondents rated search as very or extremely important, 52% of the very same respondents did not measure the effectiveness of search on their sites. The survey also indicated that only two of the 36 sites that sell online knew what percentage of people buy immediately after using search.⁶

The same firm had estimated in a previous report “that Fortune 1000 companies shell out an average of \$1.5 million to \$2.1 million per year on site redesigns without knowing whether the redesign actually made the site easier to use.”⁷

COMMON REASONS TO EVALUATE AN IA

Information architects evaluate IAs before a redesign to determine if changes are necessary and to assess what changes are needed. An evaluation can identify opportunities to improve a web site, maximizing profit and increasing productivity.

During the redesign, evaluation helps compare competing IA design ideas; for instance, whether a 5-item navigation bar performs better than a 15-item one.

After a redesign, information architects evaluate their work to demonstrate that the changes made are indeed improvements. Evaluation is often necessary to help justify the value of the investments in IA. For example:

- for justifying budget requests
- for creating information architect positions or an IA department in your company
- for hiring information architects

Decision-makers will look much more favorably upon requests that demonstrate how the IA has been improved through past investments—or, alternatively, how the current IA measures up to the competition.

CHALLENGES IN MEASURING AN IA

It is challenging to prove that the IA is the direct cause of a site’s improvement.

To prove in no uncertain terms that improvements in IA directly cause improvements in a web site’s success is extremely difficult. To do this would take a large-scale, tightly controlled study based on comparing two complete web sites—a control and a variable—that are exactly the same in every single way except that their IAs are different. Even if such a costly and difficult study succeeded for one site, it might not be generalizable beyond the site studied since the goals of users and businesses and the types of content are not identical across all sites. Practically speaking, it may not be merely difficult, but effectively impossible, to establish a causal link between improvements in IA

and improvements in site metrics. However, the converse is also true; though you can't prove there is a causal link, you also can't prove that there isn't. Note that IA isn't any harder to measure than visual design, content quality, business strategy, or most other aspects of web site development.

It is challenging to measure the IA as a whole.

Another complication is that while we can measure aspects of IA (e.g., the impact of a specific label on navigation), it is very difficult to measure the IA as whole. And the sum of all the IA parts is usually much less than the whole IA. Looking at the evaluation of an automobile design might be a handy comparison: engineers might be able to perform component testing of parts of the car, but the sum of the results to all these individual tests is not entirely the same as testing the entire car, say, by crash testing it.

Though we can't measure the IA as a whole, we can say that a particular aspect of IA is better or worse relative to another IA (e.g., a previous version) or better or worse relative to an established benchmark.

It is easier (and more fun) to evaluate a building than to evaluate that building's blueprint.

One key issue that makes the job of evaluating an IA difficult, is that, like blueprints for physical structures (e.g., a 3-bedroom house), IAs for web sites are abstract models. It is hard to see, smell, taste or touch an IA in the same way you can "experience" a web site.

In evaluating a 3-bedroom house, different types of people may ask different questions:

- A builder evaluates the house according to business goals: Will this building make enough money for me? Will this project reinforce my reputation?
- Prospective buyers evaluate the house according to overall goals: Do I like this house? Will I feel comfortable living here? Is this house a good investment?
- Inhabitants also evaluate the house according to more specific tasks: Are there enough bedrooms, bathrooms, and closet space? Is there a good place to put the mail?

Evaluating a web site is fairly similar to evaluating a building, in that a site can be measured in terms of:

- The sponsor's overall business goals.⁸
- The customer's overall goals.⁹
- The completion of a particular task.¹⁰

Note that in all of these cases, people are evaluating aspects of the house or the web site, not aspects of the blueprints or IA. And IA isn't the site; it informs the design of the site.

All these interdependencies do not mean we can't evaluate an IA; they just make it more difficult.

EVALUATING INFORMATION ARCHITECTURE: AN OVERVIEW

Site organization involves the structure, grouping, and labeling of site content. It is the bread and butter of information architecture and therefore will be the primary focus of this paper. However, many would argue that IA involves much more than organization—IA might also include areas such as navigation system design and indexing. Evaluating these other aspects of IA borrows from usability evaluation techniques, typically employing some form of heuristic evaluation. Usability evaluation techniques are well established, well documented in both books¹¹ and web sites,¹² and should be a part of every information architect's bag of tricks.

This paper focuses on evaluation techniques unique to web site organization. For evaluating a user's experience with searching and browsing, consider the following.

- Usability evaluation techniques that measure the effectiveness and efficiency of task completion, as well as user satisfaction, work well for evaluating known-item finding tasks. The ISO standard 9241-11 discusses traditional usability evaluation and measurements.¹³ Andrew Dillon has recently called for additional metrics based on: aesthetics, perceived usability, cognitive effort, perception of information shapes, acceptance level, and self-efficacy.¹⁴
- Information retrieval evaluation measures such as precision (the accuracy of search results) and recall (the completeness of search results) work well for evaluating known-item finding tasks, especially for searching. Relevance, however, is much more difficult to measure.
- SEARCH LOG ANALYSIS is an excellent way to discover where users have problems while searching. It can involve:
 - Identifying which search terms yield zero results or “too many” results.
 - Combining search logs with access logs for click-stream analysis, following the path and exit pages of users who click through a link on the search results page.

- Note, however, that known-item finding (or “hunting”) tasks do not reflect the entire range and complexity of real-world tasks. Finding, especially on the Web, tends to involve open-ended exploration (“gathering”), a shifting of goals and priorities (“berrypicking”)¹⁵ or comparisons. It also includes switching quickly between each of these different modes of finding in the same task.
- Also, finding tasks are not ends in themselves; they are the stepping-stones to larger tasks and goals. Therefore, it is best to include the end-state (e.g., the gift I ordered arrived at my Mom’s door) of the larger task or goal.

TYPES OF EVALUATIONS

IA evaluations may occur at any phase of web site development. Evaluations can measure the performance of a single IA at a specific point in time or they can:

- Compare your site to a competitor’s site
- Compare the new version to previous version(s) of your site
- Compare several different design options for the same IA design issue; for example, comparing a broad and shallow hierarchy versus a narrow and deep hierarchy for the same site

The methodologies described in this paper can be adapted for all of these variations.

Evaluation can be both quantitative and qualitative, and both approaches have their merits.¹⁶ Quantitative methodologies work better for testing a single, discrete component, (e.g., a label) or a single, discrete behavior (e.g., looking for additional “pages like this”). Qualitative methodologies are more appropriate for the squishy stuff: evaluating multiple components at once and working with users. At present, most evaluations that are being done are qualitative, not quantitative.

WHEN TO EVALUATE AN IA

The table below describes the types of evaluations that work best at the various phases of the design and development life cycle.

Phase	Types of Evaluations
Identifying requirements	<ul style="list-style-type: none"> • Obtain base-line measurements that can be used for later comparisons (Qualitative or Quantitative) • Validate existing IA designs with users (Qualitative) • Compare your IA to the IA of a competitor's site (Qualitative or Quantitative)
Research	<ul style="list-style-type: none"> • Explore users' grouping/labeling behavior (Qualitative)
Early design	<ul style="list-style-type: none"> • Validate IA designs with users (Qualitative) • Compare this IA to a previous IA for the same site (Qualitative or Quantitative)
Later design	<ul style="list-style-type: none"> • Refine/validate an IA design (Qualitative)
Post-launch	<ul style="list-style-type: none"> • Compare this IA to a previous IA for the same site (Qualitative or Quantitative) • Compare this IA to the IA of a competitor's site (Qualitative or Quantitative)

IA EVALUATION VS. WEB SITE USABILITY

Evaluating the usability of—or a user's experience with—a web site involves more than IA. Clarity of communication, accessibility, site maintenance and much more may factor into a web site's usability. However, it usually makes sense to coordinate IA evaluation experience and effort with your usability specialists. The two basic options are:

1. **Conduct your own evaluation and work with the usability specialist to adapt a usability technique so that it only applies to IA.** For example, a HEURISTIC EVALUATION might encompass visual design and technical implementation, as well as the structure of the site; however it is very possible to do a heuristic evaluation based only on IA heuristics.

Or, with PAPER PROTOTYPING—having people attempt to complete tasks using crude representations of screens, often called wireframes—a usability engineer may be looking at page layout as well as the page components.

For IAs, it will be much more effective to perform ABSTRACT PAPER PROTOTYPING—where the wireframe includes as little layout as possible—so that participants are focused largely on the path through the site organization based on the concepts, rather than being influenced by page layout. An example of how this can be done is to print out the top levels of the hierarchy—each level on a separate letter-sized page—in left-justified text using large type and a legible font and to follow standard paper prototyping techniques.¹⁷ You may also choose to randomize the order in which the items appear (on each page the participant sees) to better minimize page layout bias.

2. **Piggyback onto the usability specialist's evaluation of all aspects of the web site, but make sure that IA evaluation is included.** You might choose to add additional IA heuristics to a heuristic evaluation that addresses visual design, technical implementation, etc. Or you may choose to have the usability specialist do barebones paper prototypes (little page layout) before doing richer paper prototypes that focus on page layout.

The following three figures provide visual examples of a high-fidelity prototype, a low-fidelity prototype, and an abstract paper prototype for the Argus Center for Information Architecture (<http://argus-acia.com>).



By Argus	On Our Radar	
<p>ACIA Event</p> <p><u>Synonyms and Taxonomies</u> REGISTRATION OPEN (New York, San Francisco) A one-day workshop on thesaurus design for information architects.</p> <p>Strange Connections</p> <p>(November 6, 2000) <u>Information Architecture 2000</u> A conference trip report featuring the exceptional people and ideas of the IA2K experience.</p> <p>White Papers</p> <p>(May 2000) <u>Information Architecture of the Shopping Cart</u> by Sarah Bidigare Exploring the process-oriented information architectures of popular ecommerce sites.</p> <p><u>More by Argus</u></p>	<p>Events</p> <p><u>IA 2000</u> La Jolla, CA (10/26-10/27)</p> <p><u>ASIS Annual</u> Chicago, IL (11/13-11/16)</p> <p><u>CHI 2001</u> Seattle, WA (3/31-4/5)</p> <p><u>See all events...</u></p> <p>People</p> <p><u>Paula Thornton</u> <i>Information Architect</i> <i>Luminant</i></p> <p>Paula covers customer relationship management, Ayn Rand, metapeople, and more.</p>	<p>Content</p> <p><u>Software for Use</u> <i>By Larry Constantine and Lucy Lockwood</i> A guide to use cases and unified modeling language (UML).</p> <p><u>The Information Architecture Guide</u> A selective online bibliography of information architecture resources:</p> <p>Browse by: <u>Subject</u> <u>Author/Compiler</u> <u>Title</u></p> <p><u>Other Guides...</u></p>
<p>Subscriptions</p> <p>For bi-weekly notification of new content and links, enter your email address below.</p> <input type="text"/> <p><input type="button" value="Send"/></p>	<p>Submissions</p> <p>If you'd like to suggest events, people or content, please <u>let us know</u></p>	
[footer]		

Figure 1. Higher-fidelity view

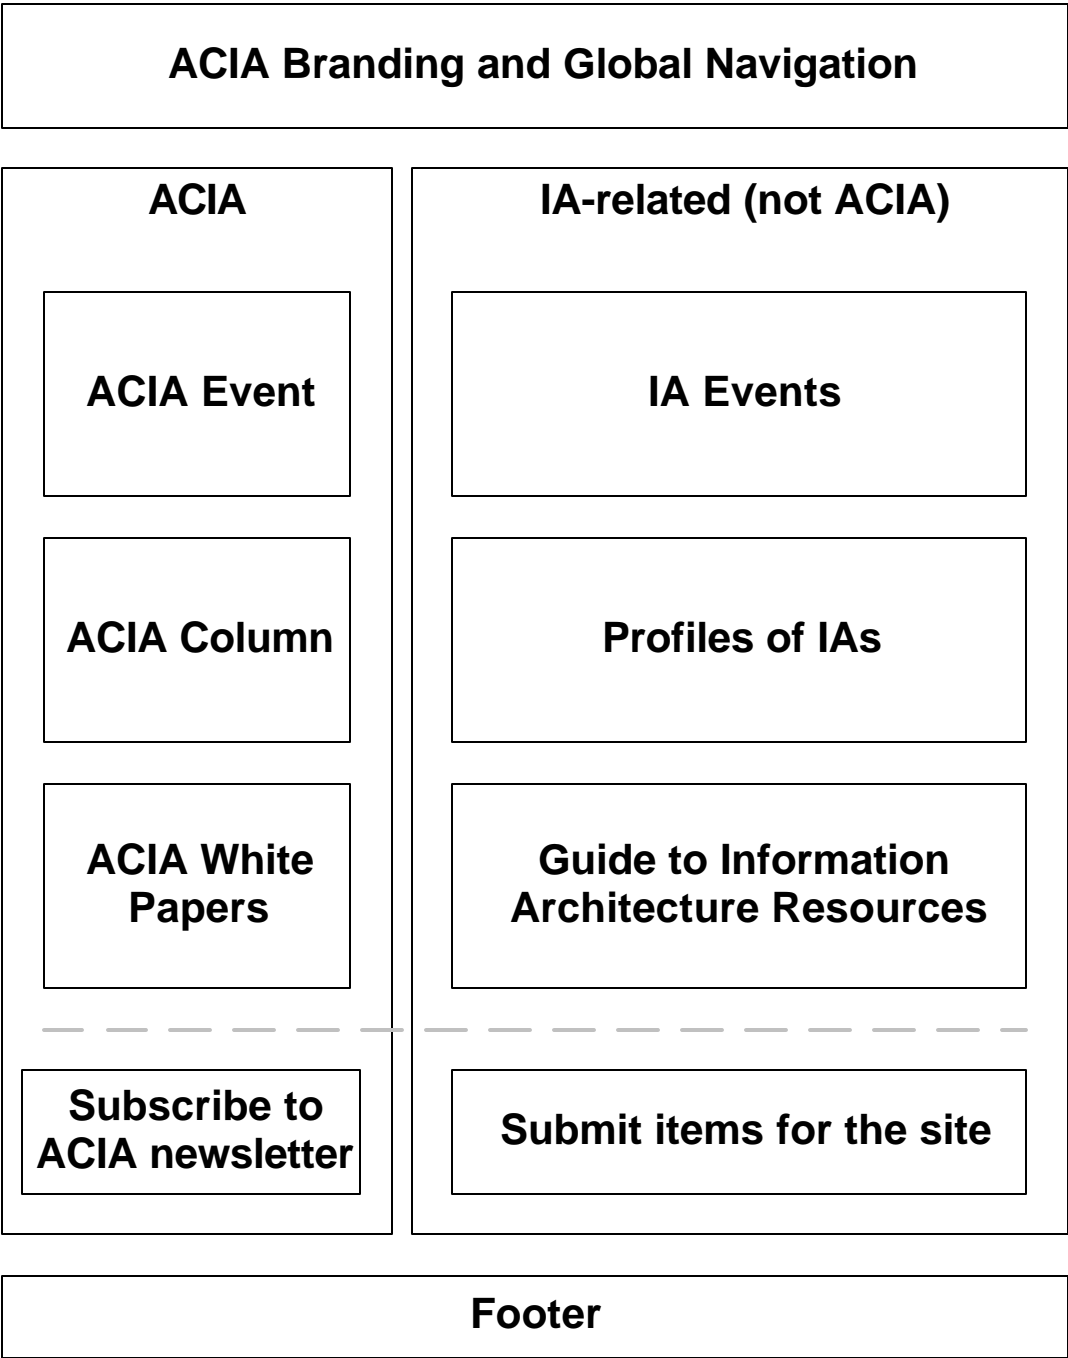


Figure 2. Lower-fidelity view

Events sponsored by ACIA

Column by ACIA's Peter Morville

Publications offered by ACIA

Events of interest to IAs

Web sites, books, articles of interest to IAs

Profiles of IAs

Request to be notified for updates to the site

Suggest events, profiles, or content for the site

Figure 3. Abstract view

EVALUATING STRUCTURE

Site structure refers to the completeness (i.e., the inclusion or exclusion of content and functionality) **and form** (e.g., hierarchy vs. database) **of the site as a whole**. Here we can measure how well the site structure reflects:

- The sponsor's needs
- Users' needs
- An appropriate structural design

We can also **measure the completeness and form of lower-level content objects, such as applications, documents, or pages**.

THE DEGREE TO WHICH SITE STRUCTURE REFLECTS SPONSOR OBJECTIVES

Description

This is one of the most fundamental of all evaluations: making sure that each of the stakeholder requirements fits somewhere in the site structure. This answers the question: Does the structure of the site meet business and stakeholder needs?

Evaluation Technique

Here is a technique for matching site structure to the sponsor's objectives:

1. Identify each user task (or piece of content or functionality) implicit or explicit in the requirements document or other documents that provide guidelines for the design of the site. If the requirements document has already pulled these together and prioritized them (by frequency, by importance to the business, etc.), this makes life much easier. If not, squeeze this step into the process.
2. Check whether or not the site structure supports all the tasks.
3. Check whether or not the site structure supports only the right tasks.
4. If the tasks have been prioritized, ensure that the site structure reflects this prioritization.

Comments

Although the process of matching the site structure to the requirements is fairly straightforward, this technique assumes that prior to evaluation, there is a requirements document that is comprehensive, sound, and has strong buy-in among stakeholders.

To Learn More

Constantine and Lockwood developed a methodology called JOINT ESSENTIAL MODELING (JEM).¹⁸ The core of JEM is a technique that allows key stakeholders to model roles and tasks and to prioritize these roles and tasks—this is a quick and effective way to define the key functionality of a site.

If you don't have a formal requirements document, a good place to start learning about requirements is *Software Requirements* by Karl E. Wiegers.¹⁹

THE DEGREE TO WHICH THE SITE STRUCTURE MEETS USER EXPECTATIONS

Description

A site's sponsor may have ideas of what users should do on the site, but users themselves often have different expectations. First and foremost, users should be able to understand the site structure, but the structure should address the tasks that users expect to be able to perform or learn about on the site.

Evaluation Techniques

Informal techniques such as monitoring user feedback work well and can be done via online forms, customer service phone calls, or comments during usability tests. It is also possible to address these issues via questionnaires (e.g., a survey that pops up in a new window after a user leaves the site) and other traditional ways of measuring customer satisfaction (e.g., surveys that measure willingness to revisit the site or willingness to recommend the site to others). Two consulting firms, Creative Good and User Interface Engineering, have been developing new techniques to listen to users and evaluate whether a site is successful based on users' goals. These techniques can be modified to focus specifically on site structure.

Comments

It is important to distinguish between attitudes about site structure and other aspects of the site such as content or visual design. For example, in one-on-one sessions with users you'll often hear them say, "This site isn't easy to use." For some people, this means they can't navigate through the site hierarchy easily but for others it might mean that the pages load slowly or that the fonts on the screen they're looking at are too small. In phrasing questions and following up, try to tease out these types of distinctions by asking "Why?" or "How?"

To Learn More

Creative Good describes their LISTENING LABS technique on pages 20-22 of their Dotcom Survival Guide.²⁰ User Interface Engineering's technique is described in recent issues of their newsletter and in the first reports of the "Designing Information-Rich Web Sites" series.²¹

STRUCTURAL DESIGN OF THE SITE AND OTHER CONTENT OBJECTS

Description

The IA design of the structure of a web site and other content objects (navigation pages, destination pages, etc.) includes such issues as the depth and breadth of a site hierarchy, the number of cross-references in a site hierarchy or an index, and the number of links from a node.

Evaluation Techniques

Pioneered by Jakob Nielsen, USABILITY INSPECTION techniques—where evaluators inspect the design—work best for evaluating structural design of a site. The most common inspection technique is a HEURISTIC EVALUATION, where several people (hopefully experts) evaluate the design, comparing it to agreed-upon design principles, or "heuristics." For example, one heuristic that your development team may agree upon is that each destination document on your site should provide at least two but no more than five links to other related destination documents and at least one link that provides access to categories of related documents. Then, using this and other heuristics, evaluators could go through the IA blueprints or other design deliverables to see if all documents meet the criteria.

Tom Tullis describes another evaluation technique that involves showing users a page mockup (low or high fidelity) or an actual web page, and asking them questions such as “Where would you click first?” for a given task or “Where would you find...?”²² This works best to evaluate both visual design and IA design at once, but you may try to make abstract mockups to minimize the influence of visual design and focus more on the IA than the aesthetics. In addition to recording the accuracy of responses, you may also record the response time: the time it takes the user to point to or say the answer.

Comments

The good news is that this is one of the cheapest and easiest techniques to use. The bad news is that there are few—if any—established heuristics or design guidelines for site structure.

To Learn More

To learn about heuristic evaluation, see Jakob Nielsen’s “Online Writings on Heuristic Evaluation” (<http://www.useit.com/papers/heuristic/>) and Keith Instone’s adaptation of Nielsen’s heuristics in the Web context (<http://webreview.com/97/10/10/usability/sidebar.html>). For more information on other usability inspection methods, see Nielsen’s “Summary of Usability Inspection Methods” (http://www.useit.com/papers/heuristic/inspection_summary.html).

EVALUATING GROUPING

Grouping refers to the coherence and integrity of categories or groups of content objects. We can measure:

- The degree of similarity of objects within a category.
- The degree of overlap between categories.
- How well a child content object represents its parent.

HOW WELL SIMILAR THINGS ARE GROUPED TOGETHER

Description

In terms of site hierarchies—tree structures with parent-child relationships—we can evaluate how well a child category or child document fits with the parent category. We can also evaluate how well documents fit with each other or how well categories fit with each other. Though it is impossible to separate the label completely from the concept, the emphasis here is on evaluating the similarity of concepts, not the effectiveness of labels (which is described later in this paper).

Evaluation Techniques

Sorting multiple concepts into predefined categories, known as a CLOSED CARD SORT (also called a VALIDATION CARD SORT, RESTRICTED CARD SORT, or P SORT) is a tried and true method for evaluating categories.

To perform a card sort, write labels—say, the titles of documents—on index cards and ask users to place each card into a category that you have defined. Closed card sorts are usually performed after the results of an OPEN CARD SORT (also called FREE CARD SORT or Q SORT), which follows the same procedure except that users create their own categories. Closed sorts can evaluate the groups that emerged out of the exploratory open sorts. It is a cheap and effective technique for determining whether a child document or category is placed in the right parent category.



Figure 4. Results of a closed card sort. Users are asked to place the white slips of paper into predefined categories: Departments and Employees, Knowledge Base, etc.

For speed, portability, and reliability, use typed labels on 3x5 index cards. However, there are also several freely available software packages that let you perform card sorts on a computer.

- Gerry Gaffney's "Classified"²³ is a Windows program that provides card sorting functionality.
- NIST's Web Category Analysis Tool (WebCAT)²⁴ provides card sorting functionality via the Web.
- IBM's EZSort²⁵ is a Windows program that provides card sorting functionality and also provides a cluster analysis of the card sort data.

Comments

Card sorting is a well-established technique in the fields of psychology and usability and is extremely helpful for evaluating hierarchies.

To Learn More

Card sorting is widely documented. One place to get started is Information & Design's "What is Card Sorting?"²⁶ To get citations to articles about Q sorts, P sorts, and a host of other techniques in the field of psychology, see Nancy J. Cook's 1994 review of "Varieties of Knowledge Elicitation Techniques."²⁷

THE OVERLAP BETWEEN GROUPS

Description

One theory states that a good hierarchy has both high within-category similarity and low between-category similarity; a hierarchy with much overlap between categories would be undesirable. This technique demonstrates a unique way to identify overlapping categories.

Evaluation Technique

To identify "problematic interactions between categories," Kirsten Ridsen tracked how users visit multiple categories while completing the same task.²⁸ The idea is that if, for example, many users visit three different top-level categories on the same task, the distinction among those three categories may be not clear.

Comments

Note that there are different philosophies about what constitutes good hierarchy design. In a true taxonomy, every item fits into only one place in the hierarchy. However, many designers choose to place an item in multiple nodes of a hierarchy under the theory that if a user is liable to look in multiple places for a single item, it should be present in all of those places. This technique addresses those who believe in the former philosophy, rather than the latter.

Ridsen suggests that a high level of redundancy makes it difficult for users to differentiate between categories. However, this does not necessarily prove that redundancy impedes effective navigation.

To Learn More

See Ridsen's paper "Toward Usable Browse Hierarchies for the Web."²⁹

Description

A common technique to clarify the meaning of a category is to list subcategory items alongside the category name. The design of the hierarchy on the Yahoo main page demonstrates this technique:



Figure 5. Yahoo main page: Directory hierarchy displays subcategory items with category label.

Highlighting sample children (e.g., Literature) alongside the parent label (e.g., Arts & Humanities) can also be very useful for providing users with context on site maps and search results.

Evaluation Techniques

Using the GOODNESS RATING technique, participants are given the names of a category and a subcategory and then are asked to rate how well the subcategory fits the category—in other words, how good is this grouping?

In a SPEEDED SENTENCE VERIFICATION technique, participants read a true-false sentence with the structure “[sub-category] is a [category].” Using the Yahoo main page, an example would be “True or False: Economics is a social science.” During these studies, make sure you capture user’s reaction times. A reasonable assumption is that quicker

responses indicate that the child is more representative of its parent; slower responses indicate the opposite.

Comments

The SPEEDED SENTENCE VERIFICATION technique applies to only a narrow problem set: which example children to display. It does not provide answers on whether or not to display children or how many children to display.

To Learn More

Karen M. Dobroth, in her article “Understanding Web Users’ Browsing Choices,” describes the two methodologies mentioned above and also introduces her own technique.³⁰

EVALUATING LABELING

Labeling refers to the name or icon of a content object, such as the title of a page, or the title of a category or heading.

PREDICTABILITY AND EFFECTIVENESS OF A LABEL

Description

Many different labels can be used to represent a single concept. Information architects need techniques to determine how well a particular label can inform users of what lies beyond that particular label.

Evaluation Techniques

Kirsten Ridsen describes a technique in which a group of participants generates labels for a category and then either they or a second group rates how well each label represents the category using a 7-point scale.³¹

A heuristic evaluation will also work well for evaluating labels. Jeff Johnson provides many good labeling heuristics (e.g., use industry-standard terms for common concepts, avoid similar concepts) in the “Textual Bloopers” chapter of *GUI Bloopers*.³²

Comments

I believe Ridsen’s technique can be quicker and more effective if the IA design team generates the labels rather than participants. Users could also generate their own label if they choose.

Although Ridsen’s technique may be time consuming if done for every label, it can be very effective for particular, problematic labels.

To Learn More

See Ridsen’s paper “Toward Usable Browse Hierarchies for the Web.”³³

ENDNOTES

-
- 1 Mark Hurst, "Holiday '99 E-Commerce: Bridging the \$6 Billion Customer Experience Gap" (Creative Good, September 1999). <http://www.creativegood.com/holiday99/>
- 2 "Online Shopping Report" (Zona Research, 1999).
<http://www.zonaresearch.com/promotion/keynote/osrprofile.htm>
- 3 Dick Stenmark, "Identifying Problems with Email-based Information Sharing," in *Proceedings of IRIS21*, ed. N. J. Buch et al. (Denmark: Department of Computer Science, Aalborg University, 1998).
<http://w3.informatik.gu.se/~dixi/publ/mail.htm>
- 4 "The Reuters Guide to Good Information Strategy" (Dow Jones Reuters Business Interactive Limited, 2000). <http://www.reuters.com/rbb/research/gis.htm>
- 5 Ben Berkowitz and Denise Levin, "IAM Sues Razorfish for Poor Design," *TheStandard.com* (July 14, 2000). <http://www.thestandard.com/article/display/0,1151,16831,00.html>
- 6 Paul R. Hagen, "Must Search Stink?" *The Forrester Report* (June 2000).
- 7 Harley Manning and John C. McCarthy, "Why Most Web Sites Fail." *The Forrester Report* (Sep. 1998).
- 8 The NetGenesis white paper on "E-Metrics: Business Metrics For The New Economy" is an excellent overview of metrics of site success for B2C sites. Available for free download from:
<http://www.netgenesis.com/emetrics>.
- 9 For information about evaluating a customer's overall goals, see the techniques of Creative Good and User Interface Engineering described later in this paper.
- 10 A good resource to learn about the evaluation of task completion is:
Jeffrey Rubin, *Handbook of Usability Testing: How to Plan, Design and Conduct Effective Tests* (New York: John Wiley & Sons, 1994).
- 11 These two books introduce a wide variety of usability evaluation techniques:
Deborah J. Mayhew, *The Usability Engineering Lifecycle: A Practitioner's Handbook For User Interface Design* (San Francisco: Morgan Kaufmann, 1999).
Jakob Nielsen, *Usability Engineering* (Boston: AP Professional, 1993).
- 12 The "Methods" section of Keith Instone's UsableWeb (<http://usableweb.com/>) points to the best Web resources on usability evaluation techniques. James Hom's Usability Methods Toolbox (<http://www.best.com/~jthom/usability/index.htm>) is another good starting point.
- 13 *Ergonomic requirements for office work with visual display terminals (VDTs). Part 11: Guidance on Usability*. ISO Standard 9241-11. (International Standards Organization, March 1998). Available for purchase from: <http://www.iso.ch/cate/d16883.html>
- 14 Dillon, Andrew. "Understanding and Evaluating the User Experience with Information Spaces." Presentation at the Information Architecture 2000 conference. Argus Center for Information Architecture, October 2000. http://argus-acia.com/acia_event/dillon_session.html
- 15 Marcia Bates, "The Design of Browsing and Berrypicking Techniques for the Online Search Interface," *Online Review* 13, no. 5 (1989): 407-431.
- 16 Peter Morville, "Information Architecture and Ulcers" (Argus Center for Information Architecture, June 21, 2000). http://argus-acia.com/strange_connections/strange002.html

¹⁷ Sharon Smith describes a similar process in her article:

Sharon Smith, "A Method for Evaluating the Organization of Content of a Web Site," *Common Ground* 8, no. 2 (May 1998): 30-33.

¹⁸ Larry Constantine and Lucy Lockwood, *Software for Use: A Practical Guide to the Models and Methods of Usage-Centered Design* (New York: Addison-Wesley, 1999): 505-7.

¹⁹ Karl E. Wieggers, *Software Requirements* (Microsoft Press, 1999).

²⁰ The "Dotcom Survival Guide" is available free from <http://www.creativegood.com/survival/>.

²¹ Jared Spool describes the methodology more briefly in his June 28, 2000 posting to the CHI-WEB mailing list. Visit the list archives at <http://www.acm.org/archives/chi-web.html> and follow this path: June 2000, week 5 → User Testing General Audience Web Sites - Advice? → the second item listed.

²² Tom Tullis, "Multiple Paths to Web Usability," conference presentation, Make It Easy 2000 (IBM, May 17, 2000). http://www-3.ibm.com/ibm/easy/eou_ext.nsf/Publish/1323

²³ Classified is available from: <http://www.infodesign.com.au/classified/default.html>

²⁴ WebCAT is available from: <http://zing.ncsl.nist.gov/webmet/>

²⁵ EZSort is available from: http://www-3.ibm.com/ibm/easy/eou_ext.nsf/Publish/410

²⁶ "What is Card Sorting?" is available from: <http://www.infodesign.com.au/usability/cardsorting.html>

²⁷ Nancy J. Cook, "Varieties of Knowledge Elicitation Techniques," *International Journal of Human-Computer Studies* 41 (1994): 801-849.

²⁸ Riden, Kristin, "Toward Usable Browse Hierarchies for the Web," in *Human-Computer Interaction: Proceedings of HCI International '99*, eds. H. Bullinger. and J. Ziegler (1999): 1098-1102. Also available online from: <http://microsoft.com/usability/UEPostings.htm>

²⁹ See note 28 above.

³⁰ Karen M. Dobroth, "Understanding Web Users' Browsing Choices," *Common Ground* 8, no. 2 (May 1998): 8-11.

³¹ See note 28 above.

³² Jeff Johnson, *GUI Bloopers: Don'ts and Do's for Software Developers and Web Designers* (San Francisco: Morgan Kaufmann, 2000).

³³ See note 28 above.

ABOUT THE AUTHOR

While at Argus Associates, a pioneering consulting firm that specializes in information architecture, Steve Toub acted as lead architect for a variety of projects. Engagements included projects with Microsoft, Applied Materials, Northwestern Mutual Life, and DaimlerChrysler. His areas of interest included usage-centered design, the intersection of IA and IT, and the intersection of IA and business strategy.

Steve Toub received his B.A. in Theology from the Georgetown University and his M.S. in Information from the University of Michigan. Steve worked at Argus from May 1997 through November 2000.

ABOUT THE ARGUS CENTER FOR INFORMATION ARCHITECTURE

Mission

The Argus Center for Information Architecture provides leadership in defining and advancing the evolving discipline of information architecture.

What We Do

The Argus Center serves as a focal point for learning about the theory and practice of information architecture. Towards this goal, we:

- Manage a selective collection of links to the most remarkable content, events, and people in our field.
- Produce original articles, white papers, conferences, and seminars that draw from the experience and expertise of the Argus team.
- Conduct research, independently and through partnerships, focused on improving our collective understanding of information architecture.

Who We Are

The Argus Center for Information Architecture was created by information architects for information architects.

It is sponsored by Argus Associates, a consulting firm that specializes in information architecture design. The entire Argus team contributes to its development.

The Argus Center also draws from the broader community of information architects, through partnerships with individuals, corporations, and universities.

Learn More

To learn more about the publications and events of the Argus Center, please visit our web site at: <http://argus-acia.com/>