In Search of Novel Ways to Design Large Cultural Web Sites

Stefano De Caro, Italian Ministry for Cultural Heritage; Nicoletta Di Blas and Luigi Spagnolo, HOC-Lab, Politecnico di Milano, Italy

Abstract

A Web site, in all domains and in cultural heritage in particular, is meant to support a variety of communication goals, like providing practical information, offering an "at a glance" understanding of what the permanent collections are about, supporting a pleasant and enticing exploration, allowing the user to locate a specific piece of content, as well as promoting the institution's brand, some selected pieces of content (shop-window effect), etc. As long as the site is small, "traditional" information architecture can cope with these needs. But when the site gets large and information-intensive, the traditional structure starts "cracking" as layers upon layers of navigation are added, and disappointment becomes a common user experience. Straight search engines have provided a reasonable solution to support just one of the above goals: allowing the user to locate a specific piece of content.

In this paper we illustrate how Rich Internet Applications (RIAs), combining lightweight information architecture with advanced search paradigms (like faceted search) and interactive visualization strategies, can be used to better support a number of communication goals. The examples are taken from the new Web site for the Directorate General of Antiquity of the Italian Ministry for Culture Heritage (to become public in Autumn 2010), where both a huge amount of content (the Italian archeological heritage) and a variety of users' profiles (from scholars to amateurs and tourists) are managed.

Keywords: information architecture, exploratory search, faceted search, Web design methodologies, rich interface applications, findability

Introduction

A cultural heritage Web site is meant to fulfill a number of sophisticated communication goals. Some of them are quite obvious; for example, offering an effective overview of the content (What is the permanent collection about?), supporting a pleasant and enticing exploration (Show me something interesting!), allowing the user to locate some specific pieces of information (Who painted Monna Lisa? What are the opening hours?), etc. There are other stakeholder goals which are less obvious but still very important, like promoting the institution's brand (e.g. "We are young and innovative") or putting forth, as in a shop-window, some selected pieces of content (e.g. the highlights section). As long as the site is small, "traditional" information architecture can cope with these needs, but when the site gets large and information-intensive, the traditional structure starts "cracking" as layers upon layers of navigation and transversal paths among them are added. Disappointment becomes a common experience for the users, who feel lost, like the visitors to the "Library of Babel":

When it was proclaimed that the Library contained all books, the first impression was one of extravagant happiness. All men felt themselves to be the masters of an intact and secret treasure. (...) As was natural, this inordinate hope was followed by an excessive depression. The certitude that some shelf in some hexagon held precious books and that these precious books were inaccessible, seemed almost intolerable.

(J. L. Borges (1941)

The search function has proven an unsatisfactory solution, as it is only capable of locating some specific piece of information, providing that the user can precisely identify it (i.e. use the same keyword the site uses), while all the other communication goals are hampered by an overloaded and strained information architecture or some extravagant communication strategies. Let us see two examples.

Example 1: searching the Louvre data base

The Louvre Web site offers access to its database of works on display: the Atlas. Let us imagine searching the Atlas for "women portrayed by women". The combination "women painter/s", gives no result. With "women paintings," three results are there: the *Death of Sardanapalu*" by Eugen Delacroix; *A Singer and a Theorbo Player Performing a Duet*, formerly known as *The Singing Lesson* by Caspar Netscher; and *Betchu and his family* (image missing), a painted limestone from ancient Egypt. Strangely enough, are there only three paintings in the Louvre somehow related to women? Should not *Mona Lisa* at least be there? A new combination, "woman portrait" gives 25 results. None of the artists is female (and by the way: *Mona Lisa* is still not there!). The advanced search is of no help. We can select the "category of work" (painting), but the other fields (like "artist") do not fit our purpose. But we know that the Louvre does display "women portrayed by women", like. for example. the portrait of Catherine, Countess Skavronsky, by Elisabeth-Louise Vigée-Le Brun.

Example 2: the MET's "featured works of art"

The Metropolitan Museum of New York's Web site puts on evidence, in the home page, a new work of art every day. An interested user is given the possibility of browsing the guided tour (next-previous) of all the featured works of art. The point is that… there are 28,196 works (information retrieved on January 25, 2010). What kind of communication goal are they fulfilling? How can the user effectively explore this huge set (let alone find something specific)?

It is clear that, in order to effectively cope with a huge amount of content on one side and the need to support a number of communication goals on the other, a new approach is required. In this paper we illustrate how Rich Internet Applications (RIAs), combining lightweight information architecture with advanced search paradigms (like faceted search) and interactive visualization strategies, can be used to better support a number of communication goals in the case of large, information intensive Web sites. None of these elements is new on its own, but the way they are designed (in view of a varied set of communication goals) and combined provides a highly effective solution. The examples are taken from the new Web site for the Directorate General of Antiquity of the Italian Ministry for Culture Heritage (to become public in autumn 2010), where both a huge amount of content (the Italian archeological heritage) and a variety of users' profiles (from scholars and professionals in the field to amateurs and "ordinary" tourists) are managed.

Background

Information architecture

According to traditional information architecture (Rosenfield and Morville, 2006), the part of a Web site that allows access to information is usually hierarchical, i. e. structured as a tree, where the root is the home page. The Web site core contents – also defined by Paolini and Bolchini (2006) as topics – (e.g. detailed information on artworks and exhibitions in a museum Web site) represent the leaves of such tree and can be "appended" to more than one "branch" (Weinberger, 2007). The topics can be in fact grouped homogeneously according to several criteria (e.g. for artworks, "all the masterpieces", "by subject", "by artist", etc.), with the aim of providing several ways for gaining access to the same pieces of content. Such "groups of topics" (Paolini and Bolchini, 2006), together with an introductive content (e.g. for "Leonardo's Masterpieces", a brief introduction to Leonardo Da Vinci's contribution to painting), constitute the access structures (the branches of the tree) to core information, and therefore are used to build up the overall navigation of the site. If the access structures are many, to reduce information overload they are grouped into one or more levels of hierarchy (this means that the outer, thinner branches are joined to thicker branches of the tree), ending up with a single taxonomical "sitemap" that encompasses the whole information architecture of the Web site. For large Web sites, however, the overall hierarchy resulting from the design process is not completely satisfactory (Crystal, 2007): users cannot easily locate what they are looking for, and interesting pieces of information are buried under levels and levels of navigation (Weinberger, 2007; Morville and Callender, 2010).

Faceted search

Search engines – both external or within the Web site – are often the only way for users to find what they are looking for. Continuing the tree metaphor, search can be considered as an automatic mechanism that "generates" the branches from a heap of leaves (Weinberger, 2007): search builds dynamic access structures (Sacco, 2006) to contents that are not pre-planned by designers and are (or should be) tailored to the specific needs of the user.

Mackinlay & Zellweger (1995) show how, already in the earlier years of the Web era, search and browsing were considered as the two faces of the same medal: navigation was in fact seen as a way for dynamically building queries on the database and exploring the results. As Web engineering and Web information retrieval developed and, in a certain sense, "diverged", such an assumption was put under discussion. Ojakaar and Spool (2001) and Spool et al. (2004) claimed that keeping users from using **search** was a best practice for usability and findability, as if search was a dangerous shortcut for designers, a sort of "diabolic temptation" they had to resist!

Indeed, a total reliance on traditional textual search (in Google's style) is far from being an optimal solution (Yee et al., 2003; Spool, 2004) for a number of reasons: the user may have a generic need, difficult to translate into a specific search query (and does not receive any good hint from the search engine); moreover, the overall communicative "message" promoted by the Web site may not be conveyed. In other words, the balance between **push** (contents that are offered by the Web site without explicit demand) and **pull** (contents accessible on demand only) would be too much moved towards **pull** (Morville, 2007).

However, in those years search was changing. New "exploratory search" (Marchionini, 2006) approaches emerged, also supported by rich interfaces (see next paragraph), transforming the search experience into a richer dialogue between the application and the user, and characterized by

iterative refinements, as in the original "berry-picking model" by Bates (1989). In particular, a better balance between push and pull can be reached with faceted search (Sacco, 2006; Tunkelang, 2009), also frequently known as faceted navigation (Yee et al., 2003; Hearst, 2009; Morville & Callender, 2010), a pattern increasingly employed for exploring collections of multimedia contents, and based on the progressive application of filters that the system combines together. By clicking on links (as in normal navigation), the user selects a combination of metadata values belonging to several classifications called **facets**. Each facet corresponds to a particular orthogonal dimension. E.g., for an artwork, there may be the following facets:

Medium: painting, sculpture...

Subject: people, landscape...

Technique: oil, watercolors...

Style: impressionism, pop-art...

Traditional Web architecture also includes multiple classifications (the "groups of topics"). The difference is that in faceted search the user is allowed to freely combine dimensions coming from different facets, thus creating *personalized* groups of topics (e.g., expressionist paintings illustrating landscapes).

Rich internet applications

"Rich internet applications" (RIAs) are Web applications with interfaces that are comparable to desktop applications, in terms of responsiveness and complexity, while in fact they are not. Different from plain xhtml pages, single elements of RIA pages may change interactively, according to users' inputs or other events, and with animation effects, without the need of (re)loading the whole page from the server.

Technologies for implementing RIAs include AJAX and Adobe (formerly Macromedia) Flash, as well as JAVA applets and other browser plug-ins.

RIA-based tools like Simile Exhibit (Huynh, Karger and Miller, 2007) can be used to implement faceted search and advanced visualization of results, even though they are currently suitable for collection of some hundreds of items only.

Web Sites and Communication Goals

A Web site is typically aimed at supporting a number of communication goals. A user may want to:

- make sense of something "at a glance"
- search for a specific piece of information
- put in relation pieces of information (painter – work of art – artistic movement – similar works of art etc.)

Moreover, a user would gladly welcome the chance to:

- stumble into unlooked-for pieces of content ("serendipitous discovery")
- create customized "groups of topics".

“Serendipitous discovery” is partially supported by the strategies like the “highlights” section (also called “director’s choice”) or the guided tours section. The second is not supported at all (being the “my museum” section a totally different – and definitely more cumbersome – way of gathering the user’s favorite items).

Eventually, there are the stakeholder’s needs. An institution, when communicating to its audience via a Web site, may want to:

- create what we may call a “shop-window” effect about its content;
- entice users to “stay and play” with the content (playful exploration);
- convey the institution’s overall brand (e.g. “we are young and innovative”) and some intended messages (e.g. “richness of content”).

Our Approach

Findability, serendipitous discovery, “at a glance” sense-making, playful exploration, branding and communication strength, as well as, of course, usability were considered critical requirements for the Web site (Figure 1). To tackle this challenge, we introduced SEE-IA (SEarch-Enhanced Information Architecture - pron. “see ya”), an integrated set of interactive and RIA-enabled design strategies that leverage existing search patterns, such as faceted search, and properly integrates them with engineered Information Architectures, to support important requirements for the user experience (see Figure 1) in large, content-intensive Web applications. The combination of search mechanisms and information architecture has already been exploited, but quite exclusively on digital libraries, archives and the like only, where information retrieval is the main user experience and therefore findability is the main requirement.

The novelty of SEE-IA lies in the fact that it blends faceted search (empowered by RIAs) with information architecture, supporting not only findability (see goals 2 and 3) but also serendipity (goal 4), “at a glance” sense making (goal 1), and playful exploration (goal 7). Strategies for properly communicating introductory content over a collection of information are also proposed (goal 5) and ,for enhancing branding and communication (goal 8).

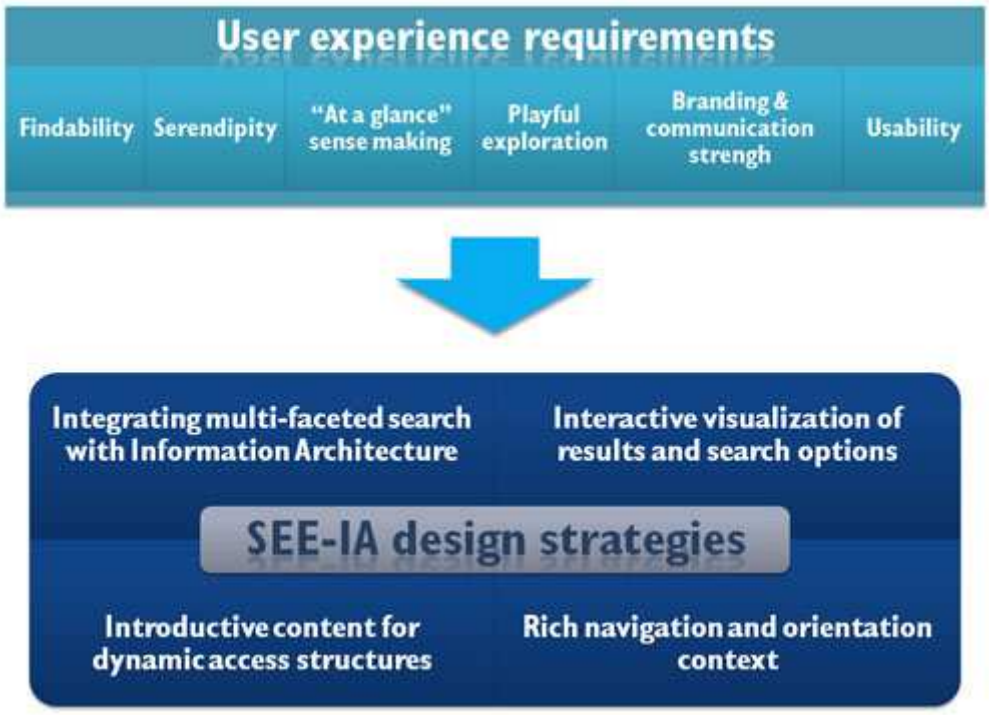

Fig 1: SEE-IA design strategies satisfy fundamental requirements of the user experience.

Integrating Multi-faceted Search with Information Architecture

The first steps of a SEE-IA design are the same as those of a "traditional" information architecture. But instead of plunging into levels and levels of hierarchy, the designer stops almost at the surface, rather concentrating on communication issues (definition of the relevant facets, as well as of the collections of homogeneous or heterogeneous items to search into, e.g. "cultural venues", including "museums" and "archaeological sites"), visualization strategies for facets (e.g. tag clouds), and search results (interactive maps).

The idea is that the hierarchy of the Web site can be simplified by designing in advance its first levels only (corresponding to the main sections) and delegating the creation and customization of the deeper levels (the group of topics) to search mechanisms. As already mentioned in the background section of this paper, faceted search permits simulating dynamic access structures. Does the user want to find museums in southern Italy related to Magna Graecia (Italian Greek colonies) civilization? No problem. Select "museums" from a "type of cultural venue" facet, "Southern Italy" from thegeographical facetand "Magna Graecia" from a cultural facet (e.g. showing the main civilizations and .periods of Italian history). If "soprintendenze" (local branches of the Ministry) are interesting too, add this value to the "cultural venue" facet, eventually getting a customized list (Figure 4).

The user can select the above facets in the any desired order, getting results after each selection in a quick, highly reactive way. These results turn out to be navigation hints that steer the interaction, like… in a dialogue! Thus playful exploration and serendipitous discovery are supported, as well as the search for something specific.

Designing interactive visualizations of results and search options

An effective visualization of both facet values and results is crucial for allowing “at a glance” sense-making. We propose to use tag clouds to visualize the facets’ values, and interactive maps and lists for the results.

Tag clouds

Values belonging to particularly relevant facets can be displayed as tag clouds, where the font size of the term is proportional to its relevance. Moreover, the size of the terms changes as interaction moves on and new selections are made (see figure 3).

Interactive maps

To enhance “at a glance” the understating and communication strength of a Web site, one option is the use of maps where results are geographically displayed. But, instead of coupling exact locations and items (that would result in a mess – were the items too many and too closely located), results are shown by means of markers, the size and colors of which are signs on their own. For example, in our case studies, three kinds of venues can be explored (museums, “soprintendenze” and archeological sites). Each of them is visualized by a marker (a circle) of a different color. Moreover, the circle’s size correspond to the number of results (which is also explicitly stated by a number in the middle of the circle itself – see figure 2). The results of the “faceted search” can be visualized in more than one way, in order to improve findability. For example, in our case study, zooming on the geographical area is possible (displaying results at region’s level, down to provinces and the exact location).

Interactive lists

Traditional lists with a sequence of items are another possibility (figure 4). The user here is allowed to sort and group items according to the same criteria of the facets.

Fig 2: Archeology in Italy. The interactive map offers the user the possibility to select the type of cultural venue (museums, soprintendenze or archeological sites: B), the geographical area (C), and the cultural dimension (D). Results are shown in the map (A) by means of circles, the color and size of which tell "at a glance" the type of venue and its relevance to the user.

Benefits and educational effects

The benefits for the user are:

- **Findability**: expert users can easily locate the venues of the type (e.g. museums), geographical area and cultural characterization (e.g. "Italics") they are looking for.
- **Serendipity**: non-expert users may discover cultural dimensions unknown to them, or unexpected locations relevant for a cultural dimension (say 'Etruscan').
- **At-a-glance sense-making**: users (whatever the level of expertise) may immediately grasp where venues (of the different types) are distributed in Italy, and their cultural characterization.
- **Branding and communication strength**: users receive a strong communication message, i.e. the richness and wide distribution of the archeological patrimony of Italy (one of the intended "brand" goals for our case study).
- **Playful discovery**: expert and non-expert users are both likely to "play" with this engaging interface to discover cultural information.

In addition, there is a remarkable educational effect: users acquire knowledge not only from predefined contents but also from something that emerges dynamically from the interaction and visualization themselves. For example, if they select the Northern area of Italy, “Romans” is the most important cultural dimension, while “Italics” is poorly represented (Figure 3-A). Selecting the Southern area, “Magna Graecia” and “Italics” emerge as relevant too (Figure 3-B). Or if they look at the cultural dimensions for Northern Italy, they may be surprised to discover that Celts are there. This is a piece of information they do not get by reading a text, but rather by playfully interacting with the application. Serendipitous “learning by doing”, which is so typical of games (Gee, 2005), is thus supported. Of course, a fundamental pre-requisite for this playful exploration is a quick and reactive interface: that is why the use of RIAs is “mandatory”.

Fig 3: Browsing Archeology (Museums and Soprintendenze) in Italy. Northern Italy (A), with Romans (“Romani”) being the most important civilization, and Southern Italy (B), where Magna Graecia (“Magna Grecia”) and Italics (“Italici”) civilizations emerge as relevant too

“Introductory content” for dynamic architectures

Groups of items; for example, “museums and archeological sites of Magna Graeciain Southern Italy”, are an important way to suggest to the user where relevant information is. A mere list of items, however, is often not sufficient. A “traditional” Information Architecture usually provides a *meaningful introduction*, by explaining; for example, what Magna Graecia was. A “traditional” search engine would instead provide a mere list of items (hopefully suitably ranked), leaving to the user the task of making sense out of it: dynamically created groups of items, such as a list of search results, may be relevant, but also “disconcerting” if not properly introduced.

Since it is obviously impossible to plan in advance an introduction specifically tailored for a group of topics that is dynamically created, we propose to associate a brief explanatory text (and image) to each facet’s value. This text can be used as a tooltip before making a selection (see Figure 4-C) and as an introductory text after the selection is made (Figure 4-B); the *combination* of the single terms’ explanations can be used as a sort of introductory text, that although not specifically tailored could still greatly help users make sense of their browsing experience. Eventually, some groups of topics could be pre-planned and therefore deserve *ad hoc* introductory texts if they emerged as relevant according to the Web site’s usage statistics or if the curators deemed them interesting.

A combination of search mechanisms and (partly pre-planned) Information Architecture will emerge over time.

Fig 4: Museums and “soprintendenze", focusing on “Magna Graecia” in Southern Italy. An interactive list (A), with introductory information (B), an interactive tooltip (C), and search history (D), is provided

“Rich” navigation context and orientation

Once users locate a set of items, say for example, “museums and archeological sites, about Magna Graecia in Southern Italy”, a number of typical actions may follow: glancing through the index of items, selecting one item and looking at its details (Figure 5), navigating to the next item (guided tour), navigating from one item to a related one (hypertext navigation), navigating back to the index for selecting another item, etc. To support these activities, “context” and “orientation” are critical. If traditional, well engineered Information Architectures are very good at this, search engines, in general, are not.

In SEE-IA, dynamically generated groups of items are “first class citizens”. They can be experienced with rich interface elements such as modal windows and consolidated navigation patterns (like indexed and guided navigation: see Bolchini &Paolini 2006) so that (i) the passage between the two types of navigation is natural and (ii) the orientation, i.e. the user awareness of the current status of navigation, is still ensured.

A search history (like the one show in figure 4-D) can be introduced to let the user go back to the previous steps of exploration, listed as links in inverse chronological order.

Dynamically generated sets can become “temporary indexes”, valid only within the current session, or can be saved, becoming a stable feature of a customized version of the Web site (available to the users who generated them). As far as links and hypertext navigation are concerned, there is no difference between the predefined set of items and the dynamically generated one. Customized Information Architecture is what we are aiming at, and what is provided by this application.

Fig 5: The result of a search: a museum’s page description. Items are opened in a new modal window that is displayed on top of the map or of the list of results

Implementation issues

After having illustrated the positive impact of SEE-IA methodology on user experience requirements, we focus here on the feasibility and reliability of implementing Web applications based on our approach.

This new generation of Web sites can be implemented using and extending reliable, existing tools, such as proprietary or open source content managements systems, as we did for our case study, where the EzPublish 4.0 open source CMS was employed, with the proper customizations. AJAX frameworks and lightweight open source tools like Simile Exhibit can be employed as good starting points for implementing faceted search and rich interactive visualizations. Search servers like Apache Solr can ensure high scalability, allowing multi-faceted searching on thousands of items contemporarily.

The high flexibility of SEE-IA makes it suitable not only for new Web applications "designed from scratch", but also for existing Web sites penalized by a too rigid and complex hierarchical organization. It is possible to apply the SEE-IA design strategy to simplify the overall hierarchy by reducing the number of levels and by reusing the metadata coming from the existing classification criteria (and additional metadata, if required) for building dynamic, multi-faceted navigation structures. RIA-based solutions on top of the redesigned information architecture will provide at-a-glance and deeper understanding, communicative impact and user engagement.

Conclusions and Future Works

In this paper we discuss the creation of a new generation of (very) large content-intensive Web sites, coupling "traditional" engineered Information Architectures (offering strong organization, powerful navigation, context orientation, etc.) with features provided by search patterns and advanced interfaces.

For the users, benefits are the possibility of easily locating what they are looking for, and most of all, the chance of engaging in a rich and educational experience where "learning" comes not only from texts, but also from the interaction itself.

For designers and developers, SEE-IA dramatically simplifies the problem of designing complex information architectures and allows them to concentrate on the communication/cultural issues directly.

Future research will consider the following aspects:

- Integration of other already existing search patterns, such as query suggestions while typing in a search box
- Adaptive combination of facet values in conjunction or disjunction, depending on the context: while in some cases combing in conjunction is desirable (e.g. "search for museums that have both Greek and Roman artifacts"), in others a disjunctive combination is more suitable (e.g. "search for museums in Italy or in Switzerland")
- Dynamic transition structures and other advanced semantic search patterns based on semantic relationships, allowing users to dynamically explore related contents (even sets): e. g. from the list of museums in Northern Italy, to the list of the Roman bronze statues in them
- Customization of the visualization tools, for example allowing the users to decide what facets they want to visualize as tag clouds
- Application of SEE-IA in other contexts, like the social Web; we are currently working on interactive filtering of discussions about a cultural "narrative" for the Cantonale Museum of Lugano.

References

Bates, M. J. (1989). "The design of browsing and berry-picking techniques for online search interface." *Online Review*, 13, 1989, 407 – 424.

Bolchini, D., F. Garzotto & P. Paolini (2007). "Branding and Communication Goals for Content-Intensive Interactive Applications". *Proceedings of 15th IEEE International Conference on Requirements Engineering*, 173–182.

Bolchini, D., F. Garzotto & P. Paolini (2008). "Value-Driven Design for 'Infosuasive' Web Applications". *Proceedings of 17th International World Wide Web Conference*, 745–754.

Bolchini, D., F. Garzotto & F. Sorce (2009). "Does Branding Need Web Usability? A Value-Oriented Empirical Study". *Proceedings of the 12th IFIP conference on Human-Computer interaction*, 652–665.

Bolchini, D. & P. Paolini (2006). Interactive Dialogue Model: a Design Technique for Multi-Channel Applications. *IEEE Transactions on Multimedia*, vol.8, no.3, 529–541.

Crystal, A. (2007). "Facets Are Fundamental: Rethinking Information Architecture Frameworks". *Technical Communication*, vol. 54, no. 1, 16-26.

Gee, J. P. (2005). "Good video games and good learning". *Phi Kappa Phi Forum, 2, 85.*

Hearst, M. A. (2009). *Search User Interfaces.* Cambridge: Cambridge University Press.

Huynh, D. F., D.R. Karger & R.C. Miller (2007). "Exhibit: Lightweight structured data publishing". *Proceedings of the 16th International WWW Conference*, 737–746.

Mackinlay, J.D., & P.T. Zellweger (1995). "Browsing vs. search: Can we find a synergy?" (panel session). *CHI '95: Proceedings of the SIGCHI conference on Human factors in computing systems*, 179-180.

Marchionini, G. (2006). "Exploratory search: from finding to understanding". *Communications of the ACM*, 49, 41 – 46.

Morville, P. (2007). *Ambient findability.* Sebastopol, CA: O'Reilly.

Morville, P. and J. Callender (2010). *Search patterns.* Sebastopol, CA: O'Reilly.

Ojakaar, E. & J.M.Spool (2001). "Getting Them to What They Want: Eight Best Practices to Get Users to the Content They Want (and to Content They Didn't Know They Wanted) ". *UIE Reports: Best Practices Series.* Bradford, MA: User Experience Engineering.

Rosenfeld, L. & P. Morville (2006). Information *architecture for the world wide Web.* Sebastopol, CA: O'Reilly.

Sacco, G.M. (2006). Some Research Results in Dynamic Taxonomy and Faceted Search Systems. *SIGIR 2006 Workshop on Faceted Search.*

Spool, J. M., C. Perfetti & D. Brittan (2004). Designing for the scent of information. *UIE Fundamentals reports.* Bradford, MA: User Experience Engineering.

Tunkelang, D. (2009). "Faceted Search". In Marchionini, G. (ed.), *Synthesis Lectures on Information Concepts, Retrieval, and Services*. San Rafael, CA: Morgan & Claypool Publishers.

Weinberger, D. (2007). *Everything is Miscellaneous: The Power of the New Digital Disorder*. New York: Times Books.

Yee, K. P., Swearingen, K., Li, K., & Hearst, M. (2003). "Faceted metadata for image search and browsing". *CHI '03: Proceedings of the SIGCHI conference on Human factors in computing systems, 401–408.*

Cite as:

De Caro, S. et al., In *Search* of Novel Ways to Design Large Cultural Web Sites. In J. Trant and D. Bearman (eds). *Museums and the Web 2010: Proceedings*. Toronto: Archives & Museum

Informatics. Published March 31, 2010. Consulted November 9, 2012.
http://www.archimuse.com/mw2010/papers/decaro/decaro.html