

Taking the Case: An Evaluation of Digital Case Study Libraries

Yeong-Tay T. Sun
Center for HCI at Virginia Tech
Blacksburg, Virginia 24060
ytsun@vt.edu

Laurian C. Vega
Center for HCI at Virginia Tech
Blacksburg, Virginia 24060
lhobby@vt.edu

D. Scott McCrickard
Center for HCI at Virginia Tech
Blacksburg, Virginia 24060
mccricks@cs.vt.edu

Abstract: Case study libraries have emerged from the development of educational resource tools as a potent contender due to their accessibility and flexibility in supporting various aspects of a learning curriculum. However, despite their advantages, the effective realization of such a tool presents a number of design challenges in its implementation. This study explores the ability of a collection of such case libraries to support its users, while revealing useful heuristics information derived from design characteristics affecting the usability of each system. Several interesting questions regarding specific deficiencies and usability issues are raised that may warrant additional exploration.

INTRODUCTION

In conceiving educational tools to better facilitate student comprehension of critical components in a design process, we must concurrently appraise user approaches to knowledge acquisition and recall in our efforts to present materials in an intuitive, imminently more learnable fashion. It has been observed that users searching for information tend to mentally structure their queries in ways that make intuitive sense to them via mental models. "Research has shown that...a good mental model [is] important to knowledge acquisition and transfer" (Ramalingam 2004), as well as maximizing the potential for realizing analogues between users' mental models and their literal data representations (Ramalingam 2004). In light of this, case studies have proven themselves to be immensely useful in the role of educational resources due to their natural inclination for presenting problems in the form of an intuitive "story-like" narrative or structure. By simulating problems in this manner, case studies effectively lend themselves to both the generation and testing of hypotheses; consequently it becomes easier to draw similarities to real-world situations and encourage logical leaps of intuition to discover parallels between previously unassociated cases (Gick & Holyoak 1980). These characteristics exemplify the underlying goals of a learning object's general-purpose instructional content by communicating information in ways that simultaneously promote key meta-cognitive skills including cognitive elaboration, error management and transfer of knowledge (Carroll & Rosson 2005). Individual cases are subject to reuse, and tend to incorporate easily into Learning Management Systems (LMS) due to their flexible composition.

Traditionally, design education has been heavily entrenched in the ideologies of observational learning, thereby surmising the effective communication of structured design principles to be an inherently difficult task. There is an abundance of issues that instructors of design principles (the core tenets of interface design) can expect to address, from choices in development methodologies such as usability engineering or goal-directed design, to specific learning objectives in the classroom. Appropriately, while an emergent focus in education centers around the development of useful learning tools to help students excel in their understanding of specific topics, it should come

as little surprise that the evolution of an educational resource to illustrate the complexities of the design process itself is fraught with such ambiguities. Methodologies, while indispensable for introducing an established curriculum from which to base course materials and exercises, are necessarily complex due to their nature and invariably retain their own deficiencies such as timeliness, lack of regular feedback and lack of breadth (Wright et al. 2005). Prior attempts at remedying these problems using case libraries have been partially successful, but satisfactory manifests of certain components in the learning object model remain elusive. In light of this, an improved tool that could successfully support the study of concepts being taught in contemporary design courses today would be immensely useful, particularly to novices of those subjects (Ardito et al. 2004). This study explores a comparison between several case study libraries, which we will examine further in the larger context of educational tools.

BACKGROUND & RELATED WORK

Obviously, a single case study would have severely limited applications in both scope and utility. To this end, we turn our attention to large collections of multiple cases, and how such a resource can be beneficial to students in an academic setting. Here, we introduce an implementation of such a library called the Usability Case Study (UCS), a collection of cases and related artifacts compiled as a learning tool for illustrating the practical evolution of a series of software projects through their respective design stages [1]. A prevailing characteristic of each of these cases is their attention to the detailed documentation of artifacts generated during each step of a project's conception. These may include requirements stages, task analysis, user-modeling, various forms of prototyping, evolutionary development, and usability evaluation methods involved in real-world instances of design examples (Carroll & Rosson 2005), ultimately allowing for the examination of cases at various levels of abstraction. The UCS was initially employed in the task of augmenting a usability engineering course's design curriculum, and had garnered a generally positive reception by the semester's end. In particular, it was reported that students, following a natural period of adaptation to the system, developed a more sophisticated understanding of usability engineering practices through their readings of homework assignments and participation in case-based activities (Carroll & Rosson 2005). The overall result appeared to be a better technical understanding of how human-centered software design works.

Aberrantly, recent experiences with the UCS aimed at mimicking the findings of this research have produced results seemingly to the contrary (Berry et al. 2006), with the much-anticipated improvements in user comprehension proving rather diminutive in scope. An informal study of the UCS's incorporation into the design curriculum of yet another usability engineering course was conducted over the better part of a regular semester. Despite the volume of materials represented, an initial complaint voiced by the system's administrator claimed that the UCS did not adequately address certain critical design implements key to helping students develop a comprehensive understanding of the design process. Allegedly, the system failed to represent the inherent complexity present in design by omitting progressional and cyclic behaviors between successive developmental phases, and by oversimplifying the represented artifacts overall. It was also concluded that user satisfaction with the system was significantly lower than expected, with multiple user reports attesting to the overly intimidating nature of the UCS towards new users, thereby lengthening the compulsory period of familiarization before producing practical results (Carroll & Rosson 2005). Others commented on the relative ineffectualness of the UCS's approach to representing a clear, hierarchical organization of artifacts within an individual case, while overly generalized summaries of inter-artifact relationships left users struggling to interpret a single artifact's relevance within the domain of a specific developmental phase. Quickly proving itself to be a hindrance for those who required further assistance, the majority of the students from this study seemed to dismiss, or were genuinely baffled by the integration of the system as part of the curriculum; that is, they were unable to grasp the relevance of this system to their current studies.

While these experiences has certainly formed a point of contention, it should be emphasized that at this time we are not attempting to make any specific assertions regarding the ability of the UCS to support its users (as there is insufficient evidence to make such a claim). The incongruity may lie with some previously undetected differences between the two groups exposed to the UCS, or it may be with the individuals overseeing the studies. However, there is also the distinct possibility that the issue concerns the implementation of the system itself; that is, a problem with the UCS may be indicative of a fundamental problem with case study libraries in general. With these results as inspiration, we wish to reassess the effectiveness of these libraries through the collection and evaluation of heuristics

[1] "Usability Case Study Library (UCS)." from <http://ucs.ist.psu.edu/>.

information to improve the design and implementation of systems similar to the UCS. That is to say, based on information collected and analyzed on our own accord, we would like some assurances that the grievances of the system's users are justified (and perhaps correlated in some manner).

MATERIALS REVIEW

Test Libraries

To this end, a brief evaluation of three different online case study libraries incorporating useful functions or distinguishing characteristics was conducted. Each of these has satisfied our criteria for a suitable implementation of a case library: they have either been integrated into the regular learning curriculum in an academic setting, and/or they have featured prominently as part of a publicly accessible website with relatively high visibility for some period of time. The focus of each library in and of itself is not especially significant, but collectively they serve as a representative sampling of actual case library implementations currently in use. Despite their differing subject matters, each website shares a number of similarities as well as unique differences, both of which we would like to examine in detail.

Pathological Case Database

<http://path.upmc.edu/cases.html>

This library comprises a fairly large database with over 480 case entries spanning the last decade or so. Photographs or diagrams of pathological data (cultures, etc.) are depicted as applicable, or appropriate. Diagnoses compare impacted organs of a patient to the healthy, non-impacted models. A wide variety of defects requiring different degrees and methods of diagnosis are presented. An interesting navigational element is incorporated by splitting each case study into two sections. The first section includes all of the details of a case, along with other relevant information and a preliminary diagnosis. At the bottom of each page is a link to the second section of the case which presents a detailed "Final Diagnosis" transcript of a patient's condition. Presumably, this separation is to encourage critical thinking by the user in speculating on their own diagnosis based on available information, rather than being accosted with the final answer immediately. This library generally suffers from the lack of powerful search and categorization tools.

UCS Similarities: Complex Search w/ advanced parsing, Mainly sidebar-driven, High-level context provided as initial view, Standard template format, Multimedia, Well-structured artifacts, Consistent formatting, Real-world examples

TED Case Studies

<http://www.american.edu/TED/class/all.htm>

This is a 14-year archive of cases (consisting of two collections bi-annually) covering trade issues affecting the environment on a global scale; many countries are included. Case studies are presented as rather lengthy articles of text; multimedia is incorporated as applicable. The large amounts of information draw parallels to real-world scenarios where it may not be immediately obvious as to what information is relevant and what is not. Each case is further divided into sections based on a unified template, while mnemonics are incorporated for rapid identification of specific cases. On closer inspection, there appears to be a great diversity among the authors/submitters for the cases included in this library. Most likely, this is a culmination of resources taken from local sources and adapted to a standard format. Although the issues of quality control and content review inevitably arise at some point, this is an excellent approach to keeping the pool of case submissions both diverse and relevant.

UCS Similarities: Key scenarios essential to case comprehension appear in their own section, along with relevant supporting materials, Standard template format, Multimedia, Artifact interpretation is open, Consistent formatting, Real-world examples

Case Law Database

<http://www.cre.gov.uk/legal/casedatabase.html>

These case studies span a number of years, with landmark cases being clearly denoted. Key points of the law that factor heavily into the subsequent ruling of a case are outlined, while a comprehensive summary of rulings appears at the end. Cases are categorized into seven sections according to a common set of groupings. The cases' contents themselves incorporate a combination of both bullet-style lists and paragraphs of text, selecting one or the other based on its suitability to the contents of a given sub-section. Summaries, by nature, tend to emphasize the notable points of an argument and would be better suited to a list-style layout. Case backgrounds, besides giving specific information, are also meant to establish a case's context and are written in paragraph form to maintain their narrative style. Besides contextual awareness, this variety in data representation discreetly separates the text into manageable chunks, resulting in an immensely more readable page.

UCS Similarities: Complex Search w/ advanced parsing, Mainly sidebar-driven, High-level context provided as initial view, Standard template format, Well-structured artifacts, Real-world examples

Discussion

You may have noticed that none of the selected case study libraries are of a usability engineering-centric, or even particularly technical (computing-wise) nature. This is mainly because the characteristics that we are attempting to gather should be those that are common among all successful case study libraries. If we were to restrict our observations to libraries that focus only on similar subject matters, we might very well be overlooking some essential feature of a successful library, perhaps one with a predisposition towards not appearing in a technically-oriented resource. Secondly, it is generally assumed that observations between similar libraries have already been conducted to some degree, and likely in the original design stages of the UCS as well. By observing these common characteristics we wish to draw conclusions from particular features firsthand based on their demonstrated worth, rather than from any particular source.

EXPERIMENTAL PROCEEDINGS

Purpose

For our purposes, what would be the most efficient means of collecting information to provide insights into the effective design and implementation of a case library? Despite the fact that such skills may only be mastered through practical experience, most people tend to learn some strategies to improve their design approaches through experience and feedback (Lewis et al. 2004). Assuming for a moment that we know nothing about designing a successful case study resource, we can speculate that the relevant points of interest will be reflected in the many interactions between a resource and its users. In this way, these points may be revealed through the careful observation of similar resources that are already widely in use, thereby enabling us to recognize the most influential and effective aspects of these libraries. In essence, this study could be considered to be a momentary step back for a more systematic assessment (to "see the forest for the trees", so to speak) before proceeding further. It is our expectation that the information gathered can eventually be used to support the development of an enhanced version of the UCS that produces a more desirable set of user responses.

Basis

Given that most case libraries are effectual only when they are both easy to use and beneficial to those who would use them, the possibilities for approaching data acquisition and retrieval are many and varied. Certain preexisting notions of optimal usability and interface layout certainly exist, but it is often difficult to discern exactly what elements result in one implementation of a learning tool being inherently more usable over another. Additionally, a virtually limitless number of factors may contribute to the success of a case study library such as user demographics, data presentation, available tools, relevance of included material, quality of content, and so on. It is evident that even subtle changes to these tools can have drastic effects on their subsequent effectiveness in communicating

information, for better or for worse. Alterations in the structure, length, emphasis, and many other components of a case can produce similar effects, and many parallels can be drawn between the two. Since we wish to explore the components of a case library that can shape a student's experience, both are equally relevant to the goals of this study, and must be equally assessed to gain a comprehensive understanding of the system.

Procedure and Overview

Our experiments involved nine (9) participants, seven (7) male and two (2) female. They were not in any inclusion or exclusion criteria; however, the intent was to seek participants who were novice designers in computer science. This was to minimize the effects of participants with past specialization in interface design or usability engineering which might skew results. Additionally, since the goal of developing this software was to help this audience design better systems, it was likewise in our best interests to select our participants from this population. Each participant was randomly assigned to a case study library before participating in a three-part experiment. The first part involved a series of tasks written on separate sheets of paper, to be completed in order. These sets of tasks, each customized to a specific use case library, were designed for probing specific aspects of the user's interaction with the system and can be found below (Tab. 1).

Question	Pathology Case Studies	Trade Environment Database	Employment Case Law
Starting from a website's front page, attempt to navigate to the case study database's Homepage, titled respectively:	"Department of Pathology Online Case Studies"	"The TED Case Studies - An Online Journal"	"Employment case law database"
Starting from a case study database's front page, attempt to locate a case study which:	is categorized under "Molecular Diagnostics": involves a man with prior medical history of Behcet's disease.	is categorized as being published in January of 1995: involves the effects of tourism on the ecosystem of national parks.	is categorized under "Direct Discrimination": involves a photographer with 22 unsuccessful applications for promotion.
Attempt to find the answers to the following questions (within each case study):	Which of the following symptoms is <i>not</i> exhibited by the patient in the initial diagnosis? [thin appearance; impaired vision; increased blinking; symmetrical reflexes]	According to the article, what island has stopped the hunt for seals?	According to the article, how many steps are there in identifying race discrimination?
Attempt to find the answers to the following questions (within each case study):	"Chemistry" data measurements in this case include which of the following? [phosphorus, barium, fluorine, glucose, potassium]	What were two factors that contributed to the decline of rug trades in Iran in the late 1980s?	What was the officially cited reason for the termination of the applicant's membership?
Identify the number of sections in this case that have appeared in other case studies explored thus far.			
From a case study database's Homepage, identify the closest means of providing user feedback. This could be an e-mail address or a scripted electronic form, for example. If none are available (or visible), navigate the site until one is found.			

Table 1: Examples of Exploratory User Tasks.

Questions ranged from tasks as simple as finding the front page of a case study library from a website's main page to locating specific pieces of information within a specific case study. The emphasis should be made that the primary goal was not to test the performance of the user, nor was it to specifically assess the correctness or incorrectness of their answers, but to develop familiarity with a given case library. Once a participant completed all of their tasks, they were asked to fill out a short likert-style questionnaire requesting feedback on various aspects of the system. The administration of these questions, which addressed issues such as system usability and overall user satisfaction, relied on the participant being able to make informed responses based on their recent experiences. Finally, an additional number of questions were administered verbally, expanding considerably upon those in part two. These

questions informally solicited opinions such as the sort of activities one might use a given library for, or how certain characteristics of the case studies influenced the participant's impression of the library. We generally considered this last portion of the experiment to provide the most useful information since we could gauge the participant's opinions of the system directly.

RESULTS AND ANALYSIS

A goal of any worthwhile educational tool is the ability to retrieve specific information on demand without asking users to inordinately burden themselves with the task of data acquisition. On an intuitive level of reasoning, one could allegorize a person's "natural" approach to seeking data as the sequential parsing of text -- that is, traversing a well-ordered physical entity implicitly defining the flow of information for some specified volume of data. However, scrutinization of participant behaviors in our study indicates that users tend to sift through textual data in one of three ways: "reading", "scanning", or "searching", in order of decreasing granularity of text. Areas of lower density will be methodically "read" in greater detail, whereas areas of higher density will be "searched" through in a concentrated effort to reduce the amount of information to be parsed. Evolving a case's foundational level of complexity beyond specific (though independently variable) thresholds compels users to scale their search patterns accordingly. While the need for a unified search methodology is intuitive and expected for such activities, a preoccupation with data acquisition suggests that the current techniques used for volume management may be inadequate for sufficiently minimizing the levels of indirection in propagating users towards their desired data, thereby prolonging and convoluting the learning process overall. From the averaged results of the electronic survey from step 2, we saw that participants overwhelmingly identified the third case library (Employment Case Law) as yielding the most satisfactory experience, while the second case library (Trade Environment Database) garnered the lowest average score. Not surprisingly, these libraries represented those with the shortest and longest average case lengths, respectively. The fact that users preferred the shorter cases by far identifies a sort of learning deficiency associated with longer cases, not necessarily on virtue of their content, but from intermediary concerns such as content management and presentation of data. This suggests that information density does, indeed, play a significant role in the perceived usability of a case and has specific implications for an instructor's approach to structuring a case study, whose contents are often derived from comparatively unabridged materials such as textbooks.

Surprisingly, we observed many participants refraining from using the available search functions at their disposal during certain stages of data acquisition. Rather, they preferred to scan the presented information visually, and in sequential order from top to bottom until this scanning approach was deemed to be infeasible. This behavior was exhibited by one participant who, after locating the relevant information for a task through linked teleportation, subsequently returned to the case's artifact hierarchy overview to perform a manual traversal of the data. When asked why he responded, "I want to see where [the information] sits [in the artifact]." In the scope of learning tools, this further reinforces the notion that multiple views of information from different spatial and organizational perspectives can help users to better comprehend the significance of that entity in the larger picture. A common complaint with overly dense case libraries targeted the overall hierarchy of information as being convoluted to the point of being unusable, with specific instances of users even becoming "lost" in a library because they were no longer aware of their relative position within it. In another task, a user was apparently unaware that two similarly presented pieces of information were actually located on the same page. This demonstrable loss of relative hierarchical referencing, or context, from within a case is considerably problematic since they can hinder tasks that are vital to supporting information retrieval. For example, a loss of context here effectively interferes with a user's predisposition towards gathering supporting information for the purposes of data validation. Logically, the greater the number of supporting materials, the greater the chance that an answer would prove to be useful to the user. In the absence of context, the ensuing search for supporting materials is greatly compromised, and any useful information to be gained due to spatial locality is noticeably attenuated.

Interestingly, this loss of context appears to be an indirect consequence of the increasing trend of content standardization across cases. Given the physical similarities in their representation, it is paramount that some mechanism be in place to clearly distinguish between successive artifacts, and hints at the need for navigational tools in a case library to provide contextual information: for example, where a user logically resides within the overall hierarchy of a given library. Within text-heavy environments, we did notice a user's tendency to visually latch on to any distinctive-looking features of a page. In particular, participant feedback revealed that the mere

inclusion of multimedia elements such as pictures seemed to have a marked effect on the perceived usability of a case, regardless of where the picture was positioned or how it was organized within the case body. When asked, participants said that in a way, the use of pictures allowed them to gain a high level understanding of what the case was about. Accordingly, it also made participants feel that the authors of the case study library had taken more time in crafting that particular case (thereby justifying their reaction to that case being somehow superior to the others). These responses may reflect the role of those images as the symbolic representation of clearly identifiable "landmarks" in the case body. In spite of the rather arbitrary usage of graphical elements seen here, such landmarks can be very effective as section/subsection delimiters or subtle content markers without the need for explicitly designating the content as such. Using the principles of spatial locality, such landmarks would presumably go far in guiding a user's attention to adjacent areas of discussion and improving information recall within a given case. Ultimately, the degree to which such actions are supported throughout the library may directly affect a student's ability to comprehend the presented information.

Together, these observations illustrate a number of issues stemming from cases that fail to adequately balance density, length, organizational structure, or similar issues thereof. While an idealized user behavior should be both predictable and consistent, the results of this study suggested that most participants did not have a clearly definable plan for escalating their searching and navigation patterns in the face of unexpected deterrences, instead falling back to some sporadic cycle of the three basic parsing methods mentioned earlier as a coping mechanism of sorts. Clearly, there is some incongruity between the behavior typified by users interacting within the system and the previously defined "natural" approach to parsing data. While the initial lack of user planning can be attributed to a predictable unfamiliarity with a system, further navigational complications implicate the absence of a strong spatially-oriented sense of organization within the library. Given that an ideal learning object should assume certain properties of a self-guided independent study and is consequently dependent upon an strong, underlying delineation of materials, this is a considerable cause for concern as it will also prove detrimental for data representations that revolve around related properties (such as temporally-arranged timelines). This equates to a need for additional constructs to help guide a user's interactions, particularly in the initial stages of familiarization. Feedback sessions revealed that many users were well-aware of these deficiencies in information management, and alluded to the importance of being able to "drill-down" a piece of content to simplify or expose data to the desired level of complexity. Being able to control the presentation of data allows users to disseminate materials at their own pace, since not everyone learns at the same rate. As these examples demonstrate, perhaps the hierarchical organization of information could also benefit from an application of Schneiderman's mantra of "overview first, zoom and filter, and then details on demand" (Schneiderman 1998).

Previous attempts at improving the UCS yielded the visualization tool VaDeR (Berry et al. 2006), which sought to minimize the degree of abstraction between artifacts by taking advantage of the principles of locality and icon-based encapsulation to eliminate the need for a prominent navigational component. As we can now see, a notable deficiency in VaDeR's implementation was the decision to eliminate the hierarchical organization of information for the sake of navigational simplicity. This action effectively aggregated all of the interface components into a close proximity, resulting in an unwieldy presentation of interface components. Additionally, the underlying logical organization of these elements was not as strongly mirrored in its physical implementation, nor was the interface particularly intuitive to use. The persisting challenge lies in structuring materials such that users can locate the data that they seek without inadvertently overlooking it due to their instinctive information parsing sensibilities.

CONCLUSION

This paper has presented a brief introduction to the appeal and usefulness of case study libraries in academic settings, as well as motivations for exploring some of the pitfalls and constructs necessary for a more effective implementation. The Case Law library appears to have garnered the most positive impressions from participants by far, thanks to its concise narrative, consistent presentation of information and uncomplicated language. The "glanceability" of this library was also improved by the subdivision of the main body of text into a series of easily accessible sections. The pathology library elicited the most neutral responses from experimental participants. It's use cases generally possessed the brevity shared with the Case Law library, but were hampered by an abundance of technical jargon, with a corresponding effect on participant performance. Finally, despite the extensive use of multimedia, the TED library cases proved to be the least popular to participants due to their incredible length and extensive supporting materials, marginal search capabilities for a database consisting of well over five hundred

cases, and a veritable gale of perpetually shifting page layouts, formats, colors, and virtually everything else from case to case.

FUTURE WORKS

Online educational resources are in a continual state of evolution, and from the results of this study we wish to compile a set of heuristics that will be indispensable in our future assessments of similar resources, including case libraries. Ultimately, we wish to build upon the UCS to a point where its functionality approaches the results found in the original trial study, and will serve as a useful learning reference for students in their study of the principles of design. The prototyping of an improved case study library centered around a timeline-based view of development is currently underway that will incorporate elements of this work in its final implementation.

REFERENCES

Ardito, C., D. Marsico, M., Lanzilotti, R., Levialdi, S., Roselli, T., Rossano, V., & Tersigni, M. (2004). Usability of E-learning Tools. Proceedings of the Working Conference on Advanced visual interfaces. Gallipoli, Italy, ACM Press: 80-84.

Berry, B., Hobby, L., & McCrickard, D. S. (2006). Exploring Benefits of Visualization for Case Studies. World Conference on Educational Multimedia, Hypermedia and Telecommunications 2006. P. Kommers and G. Richards. Orlando, FL, USA, AACE: 485-492.

Carroll, J. M. & Rosson, M. B. (2005). "A case library for teaching usability engineering: Design rationale, development, and classroom experience." *Journal of Educational Resources in Computing* 5(1): 3.

Gick, M. L. & Holyoak, K. J. (1980). "Analogical problem solving." *Cognitive Psychology* 12: 306-355.

Harrison, S., Back, M., & Tatar, D. (2006). "It's Just a Method!": A Pedagogical Experiment in Interdisciplinary Design. Proceedings of the 6th ACM conference on Designing Interactive Systems. University Park, PA, USA, ACM Press: 261-270.

Lewis, T. L., Rosson, M. B., & Pérez-Quñones, M. A. (2004). What Do The Experts Say?: Teaching Introductory Design from an Expert's Perspective. Proceedings of the 35th SIGCSE Technical Symposium on Computer Science Education. Norfolk, Virginia, USA, ACM Press: 296-300.

Schneiderman, B. (1998). *Designing the User Interface: Strategies for Effective Human-Computer Interaction*. Reading, MA, Addison Wesley.

Wright, T., Noble, J., & Marshall, S. (2005). Using a System of Tutorials and Groups to Increase Feedback and Teach User Interface Design. Proceedings of the 7th Australasian Conference on Computing Education. Newcastle, New South Wales, Australia. 42: 187-192.

Ramalingam, V., LaBelle, D., & Wiedenbeck, S. (2004). "Self-Efficacy and Mental Models in Learning to Program." *ACM SIGCSE Bulletin* 36(3): 171-175.