

Technologies of Information: HCI and the digital library

Andrew Dillon

Indiana University

adillon@indiana.edu

This item is not the definitive copy. Please use the following citation when referencing this material: Dillon, A. (2002) HCI and the technologies of information. In: J. Carroll (ed.) HCI and the Millennium: New York: ACM Press/Addison Wesley, 457-474.

Abstract

Digital libraries represent the latest and perhaps the most important stage in the transformation of our relationship to information and its supporting technologies. While the World Wide Web has raised to broad consciousness the ideas of rapid, remote access to vast collections of distributed, hyperlinked documents, users are as often confused and disoriented by such resources as they are empowered. As we enter the new millennium the possibilities for new forms of information seem boundless. Meeting the design challenges requires HCI to offer valid, reliable and timely guidance. At the same time, the creation of digital libraries provides a research forum for HCI that is rich, relevant and receptive to our inputs. The present chapter will briefly examine the initiation of the digital library concept, the early HCI involvement in this domain and present a vision of the future of HCI research and design for new information technologies.

Introduction

Though not precisely defined in current usage, the term 'digital library' is frequently employed to describe any organized and networked collection of information that is stored, accessed, and presented electronically. The term commonly refers to large repositories incorporating some combination of text, graphics, animation, and even video, accessible either locally or over the Internet, so we might reasonably conclude the size

and mixed media are further elements to throw into the definition pot. But these are largely computing requirements. As Fox et al (1995) note, the term 'digital library' connotes various images: the digitization of existing libraries, a collection of distributed information services, networked multimedia or even collaborative work environments, depending on who is imagining it. Thus, DLs occupy and indeed create a social space for users and producers of information, the design of which raises significant human factors issues that lie at the core of HCI's contribution to knowledge.

To understand the excitement of Digital Libraries and the interest they hold for the field of HCI one needs only to consider what users can or might do with such resources. No longer dependent on the physical artifacts of information such as books, journals, pictures etc. users might access from their immediate environment information that only exists in another country. Rapid reproduction of the information enables copies to be made instantly, perhaps for little or no cost, and subsequently modified, marked-up or dissected to suit the needs and preferences of the user (as opposed to the intentions of the author or the requirements of the publishers). Beyond access, users can exploit the technology to search, analyze and extract patterns from the information before reporting the results of such analyses instantaneously to a worldwide audience of interested fellow users. Multimedia configurations can enable modeling and simulations explorable by learners on demand. Filtering software might tailor both delivery and format of information for you, regardless of your point of access. In short, Digital Libraries hold forth the promise that information can be stored, accessed, manipulated and moved in ways and at speeds that were heretofore impossible. The implications for designers and users of information technology are profound. Indeed Marchionini (1999) considers the term 'Digital Library' to be too constraining for what is on offer, and suggests that 'Sharium' as a better conveyer of the potential for such technologies to create shared workspaces for information exploration and collaborative work. For the remainder of this chapter it is worth keeping such visions in mind whenever the more standard term digital library occurs.

The present chapter is concerned with the application and impact of HCI theories and methods on the design, use and acceptance of Digital Libraries. It is not a review of

existing digital libraries nor even an explanation of how digital libraries work. Instead it is an examination of how HCI can contribute to the development of better information technologies, with DLs as the focus of my examples. Also, since we are discussing HCI, my treatment of DLs is user-focused, not technological, which places this work in sharp relief to most of the DL literature that has emerged to date. As I will attempt to show, the emergence of DL systems is both a boon and a threat to the field. The technology's importance makes HCI central, while the nature and presentation of HCI as a discipline threatens to sideline our findings and methods. My intention is to point to a future for HCI that maximizes the boon side while overcoming (not just minimizing) the threats.

While digital libraries are at the forefront of technological advances, the HCI concerns are more familiar. Indeed, there are many parallels between the role of HCI in this domain and the issues faced 20 years ago by HCI researchers seeking to shape and understand the impact of desktop computing by non-specialist users. I will return to this theme later as I look ahead to HCI in the 21st century, but for now, we need to establish a context. For digital libraries, that means reflecting back briefly to the days before digital computing, and the roots of the shift from paper to alternative information technologies.

Antecedents of Digital Libraries: the ideas and the evidence

The major thinkers

The article most often cited when discussing the origins of the digital library is Vannevar Bush's *As We May Think* (Bush, 1945). Bush has been described by many as a visionary who foresaw the increasing specialization of knowledge workers and a commensurate need to access an ever growing collection of information. For Bush, a means of representing and accessing this information was essential if knowledge was to be organized and nurtured. He presented a conceptual design of a system to support the location, use and addition of information which he termed the 'memex' - a mechanized workstation that would enable retrieval and linking of information.

There are innumerable references and accounts of Bush's work in the literature on Digital Libraries (see e.g., McKnight et al 1991, Nielsen, 1995) though nothing really improves

on reading the paper itself. I shall not recount more of Bush's memex here except to point out that the idea of linking information based on personal meaning to form "trails" was a striking idea which is now accepted as standard by Web users. Indeed as Lesk (1997) notes, for 40 years, Bush's ideas of personal associative linking seemed out of fashion in a worldview dominated by the application of statistical analyses of language to computer-based information retrieval. The emergence of the web and the demonstrable willingness of users to exploit its linking functions shows how attractive personal association is to the user community, vindicating Bush's original argument.

While the conception of the digital library is attributed widely to Bush, he really was not thinking in digital terms. Others more directly saw the computer as the leverage for new forms of information technology, among them Ted Nelson and Doug Engelbart. Nelson is credited with coining the term "hypertext" to refer to the linking and direct access of information nodes, foreseeing the possibility of networks of information, accessible from any computer, incorporating everything ever written. He sought the creation of a docuverse, an information structure in which the entire literature of the world could be linked (Nelson, 1988). Calling his dream system "Xanadu", Nelson argued that nothing ever needed to be written twice since a document could be built up of original bytes and links from other documents. In many details, Nelson clearly envisaged a worldwide network that has many parallels to our current technology.

Engelbart's emphasis has always been on augmenting or amplifying human intellect. He envisaged technology as a means of empowering users to perform tasks that would prove difficult or impossible for them to perform on their own. From the DL perspective Engelbart is key since he envisages knowledge workers interacting online, with the computer providing the means of storing, accessing and presenting a wealth of information on demand. What makes him unique is his recognition that the exploitation of such a technology might require training or new skill acquisition on the part of the user.

Writing more directly about hypertext than digital libraries, McKnight et al (1991) argued that we can see Bush, Nelson and Engelbart as representing three different views of information technology which continue to attract adherents today. Bush advocates

information storage and presentation through associative mechanisms that reflect in some sense the underlying structure of the human mind. Nelson most directly seems to have envisaged the Web, with users able to access any document from their own workstations. Engelbart seeks an augmentation environment; the users of information technologies should be able to achieve more through using information technology, even if they required some learning or training.

Each of these views resonates with HCI concerns such as providing direct access, building tools that are user-centered and cognitively compatible. However, these ideas could not be directly tested until the technology emerged which more fully embodied them. The interaction underlying Digital Library design and use only began to emerge once the technology took a form that was amenable to user testing (a common criticism of the traditional HCI perspective on design also, that it routinely becomes involved only after a design has been built). For Digital Libraries, that technology was hypermedia (a term that I use generically here to include hypertext and multimedia, the technology which enabled the Web) - the emergence of applications that enabled the active linking and direct accessing of information nodes within a larger document or document set.

HCI enters the Digital Library.

Hypermedia technologies emerged dramatically in the late 1980s amid claims for the death of the book, the advancement of interactive instructional technologies and the emergence of a new digital age. I have written elsewhere about the precise myths underlying hypermedia use (Dillon, 1996) but in a nutshell these include arguments that:

information could be modeled on the cognitive structures of human memory;

users would be liberated from the tyranny of paper's linear forms;

rapid access to information would automatically lead to better use and/or learning from hypermedia than from paper;

future gains in speed and storage would render any current problems obsolete (the classic deterministic argument so pervasive in utopian views of information technology).

HCI researchers immediately began to test many of these emerging applications. Over the next 10 years a stream of empirical investigations compared hypermedia with paper documents, tested various alternative hypermedia formats, and analyzed user performance with numerous hypermedia features. Summing up these findings in 1994 I concluded that the experimental findings centered on five basic outcome differences (Speed; Accuracy; Comprehension; Fatigue; and Preference) and two major process differences (Navigation and Manipulation). It became clear from these experiments that the claims made for the superiority of hypermedia technologies for information tasks were frequently inaccurate. In truth, users, depending on their tasks, frequently experienced tremendous difficulties with this technology and routinely performed faster or more accurately with paper (see Dillon, 1994 or Muter 1996 for detailed reviews of important human factors findings on user performance with digital documents).

Just as it might have seemed that HCI was beginning to understand some of the important variables in digital information design and would be able to impart recommendations to designers, hypermedia creation and use exploded with the emergence of the World Wide Web. With an uptake rate that outpaced all expectations, Nelson's 'docuverse' was given form, and Bush's 'trails' began to emerge. What might have been once viewed as a standalone technology became a worldwide network, and the user response indicated a huge demand for the tools to create, search and link large dispersed multimedia resources.

As we enter the 21st century, the interest in Digital Libraries is reflected in major international funding initiatives, well attended annual conferences worldwide, and the emergence of dedicated academic journals on the subject. Digital libraries are being created across the globe, covering subject domains as varied as musical composition (e.g. VARIATIONS at Indiana University), to the history of the US (e.g., the Making of America Project jointly developed by Cornell University and the University of Michigan). But such examples constrain as much as they illuminate since a simple search

of digital library on the web will likely result in more output than any user could usefully exploit on topics one barely believed existed.

From an HCI perspective, much of the literature on DL research is highly technocentric. Papers report on how a DL was constructed, how image files were treated, or how knowledge-based metadata can be extracted, using specific projects as examples. Users only occasionally raise their heads in such a literature, for example, being studied in less than a third of the papers published in the latest ACM Digital Library Conference Proceedings (DL00). However, such a statistic can be misleading. The power of this technology imparts concerns with ownership, copyright, privacy, accessibility, authenticity, transfer and permanence - issues that extend beyond the technical to the psychological, social, economic and legal domains. Literatures are growing up around each of these topics so rapidly that one can hardly hope to gain and keep perspective through a single journal or conference proceedings. Though the labels might be different, concern with digital information pervades a range of conferences from educational technology to geographic information systems, and from medical informatics to literary theory. Among the issues raised by DLs, usability and the human response to digital information, remain key, and for the remainder of this chapter, it is the contribution of HCI theories and methods to DL design, development and use that will be my focus.

HCI research: from enabling to envisioning

In reviewing the last 20 years of HCI research on information design, its application to digital library design, and in speculating where current initiatives in might take us, it seems to me that the field of HCI can be dissected into three somewhat distinct research emphases. For present purpose let us call them stages of research, each embodying a set of assumptions, emphases, and methods about what questions are most important to answer, and how we might proceed to answer them. In the sense that I am speculating on HCI in the 21st Century, stage three is clearly less developed, although the future seems to call for answers to some of the first questions raised by the earliest thinkers in Digital Libraries. In the present section I will outline these stages and discuss their various

emphases and contributions to shaping our understanding of interaction and HCI's creation of knowledge.

Stage 1 - Interface design and the methodological tradition

What quite rightly identifies HCI as a distinct emphasis in systems design is its strong empirical tradition forged in the earliest work of people like Shackel (1959), Shneiderman (1980) and others who advocated user testing as the best means of improving the usability of the human-computer interface. Since the 1970s, HCI researchers have routinely run experimental trials of interface features and published the results in mainstream HCI journals and conference proceedings, creating en route a resource for designers and other researchers seeking answers to questions on such topics as optimum menu depth, screen size, input device, response rate, image quality etc.. Like the classic human engineering tradition from which it sprang, this emphasis embodies the experimental paradigm and research methods of psychology and related behavioral sciences, treating interface features as experimental variables to be manipulated and their impact measured. Such work provides the bedrock for better design and DL research is still sadly lacking in such formal user studies (though see Shneiderman et al 2000).

While feature comparisons have often proved illuminating, a research stream focused primarily on features is always going to be partial in its coverage, and in the mid 1980s, a move to couple the approach of user testing to the iterative design of new systems led to the emergence of usability engineering (Gould and Lewis, 1985). Advocating an operational definition of usability (usually incorporating effectiveness, efficiency and satisfaction as dependent variables), usability engineers set measurable criteria for user performance that any interface must support. User testing is advocated as early and as frequently in the design process as is required to either confirm performance goals are met, to provide clues as to how the interface might be re-designed to meet them, or as happens more often, to confirm that the appropriate scenarios of use are being envisaged and supported.

Usability engineering has proved extremely powerful as a means of bringing HCI into the design process early on, a classic problem for the field. Yet it would not be too harsh to describe the approach as fundamentally empirical in nature, eschewing theory in favor of performance data. As such, usability engineering is concerned less with *why* an interface design works than with demonstrating that it *does*. Landauer (1991 and 1995) probably is the most articulate proponent of this methodological approach to HCI and he makes an eloquent case for why HCI should pursue the empirical route. While the pragmatics of his case are seductive, it is my view that we cannot be atheoretical as a field, even if we so wished, since the very artifacts we design must embody theoretical assumptions no matter how weakly articulated (see e.g., Carroll, 2000, Dillon, 1995). Failing to attend to these theoretical assumptions weakens the initial prototype and renders interpretation of test findings problematic. Furthermore, HCI has always sought be more than just about interface design, but about interaction. It is against this backdrop that stage 2 research in HCI has gained a foothold.

Stage 2 - Modeling interaction - the theoretical tradition

It is not yet possible to talk of a complete theory of interaction, and many believe we will never have such an all-embracing theory. However, as findings on user performance accumulated, attempts have been made in the HCI community to broaden and generalize these data, interpreting them in the light of theoretical positions, usually incorporated from psychology, but also in the 1980s and early 1990s from sociology, anthropology, and related fields.

While part of this approach has been driven by concerns about collaboration and distributed cognition and the potential limitations of a narrow cognitivist approach, theoretical efforts also sprung directly from interface analysts seeking generalized principles and laws of interaction. This approach has worked best where it has been constrained to explain specific interactive phenomena rather than the full range of user responses to information technology. Space precludes a detailed examination the range of HCI theories here, but examples of specific phenomena that are theoretically explicable include the reading speed differences for digital and paper media (Gould et al, 1987),

expert performance speeds for routine (non-decision making) cognitive tasks (Card et al 1983), and user performance in menu navigation tasks (Norman, 1992). Even the supposedly complicated dynamics of user acceptance of new technologies show signs of being partially predictable with current models (see e.g., Davis et al 1989).

While none of these theoretical explanations constitutes more than a localized model of some facet of human-computer interactions, their emergence and application represents a level of maturation in HCI that should not be dismissed. Of course I am emphasizing theoretical positions that both predict and explain interactive phenomena. There exists a multitude of theoretical positions that can be brought to bear, each offering a plausible explanation for some form of human-computer interaction. But it is the more thorny issue of prediction that separates the useful from the potentially informative or the post-hoc rationalization. We are a long way from having sufficient theoretical power to predict many of the user issues that are important to usable systems design, but we are no longer completely dependent on user testing to determine the design alternatives we consider. All of the examples listed above could usefully be applied to some aspect of DL design and implementation. Furthermore, weaker, but no less valuable predictive power can be derived from case studies and involvement in design processes, whereby the importance of HCI professionals roles as facilitators, early testers, scenario-generators and task analysts can be confidently established and used as estimators of a successful design outcome. It is not difficult to imagine that 21st Century HCI will extend and improve its theoretical base to cover other facets of interaction. In so doing, the move from being an evaluative to a prescriptive discipline should provide a push for HCI to move on to stage three.

Stage 3 - Beyond usability: enhancement and the design of augmenting technologies.

While the field continues to seek theoretical explanations of greater power and generalization, there exists a third level of emphasis, one that the field has not yet truly attained but which remains a goal. I will call this stage 'enhancement'. By enhancement I mean to describe HCI's ability to lead the design of technologies that truly empower users, to support them in the performance of tasks that would be impossible otherwise, or

to enable users to overcome limitations in their own capabilities that hinder their development. The term is a very deliberate reference to Engelbart's call for augmentation, which has largely been ignored in the years since he first advocated it by HCI's resultant focus on usability.

While I do not wish to play down the importance of usability, it strikes me that on its own, usability is really not a sufficiently ambitious goal for this field. If it were, we should be content to remain an interface design speciality expending effort either refining our methodological tools to speed up the collection of reliable and valid data on interactions or devising more localized models of tasks. Certainly this is not a bad goal, but determining and measuring usability is neither representative of everything HCI researchers wish to do, nor does it encapsulate all that I believe the field can offer. In a context where technology permeates both work and leisure activities, traditional notions of effectiveness and efficiency are rightly questioned as potentially limiting.

As it evolved, HCI sought ways of studying human performance and working practices at a deep level, trying to understand how tools are embedded more in our collective practices, and how technology might empower users to attain performance and knowledge gains that are beyond their reach without technical support. While the plurality of approaches at work here makes it appear to a casual observer that perhaps such work is either too diffuse, or is not really HCI, I would counter that the analysis of interaction can not be so neatly bounded to interface concerns, especially if we are to serve HCI's primary aim of supporting the design of more humanly acceptable information technologies.

To grow, the field needs to be more predictive than evaluative in its scope, capable of contributing to the identification and analysis of usage scenarios where new technological forms might be envisaged to enhance human capabilities in all spheres. At this time, it is clear that enhancement is not where HCI has made its major contribution and it is difficult to point to examples of research that embody this level of issue. I remain hopeful, however, that this is the way forward for the field, and writers reviewing HCI, if it exists at the end of the next century, are likely to see enhancement as a defining

characteristic of such a field. What is exciting about digital library technologies, is the power they offer ordinary users (whom I believe will be the primary drivers of information technology designs in the next century), and it is this power that I believe can promote the study of HCI to a new stage.

Problems with HCI's role in Digital Library design

While I have avoided placing too rigid a timeline on the three stages of research, it should be apparent that they have followed a relatively sequential path, from interface analysis, to theoretical models towards, hopefully, the enhancement stage, with large overlaps. Though, worryingly, each new technology seems to induce a recycling of these stages with minimal attention given to cumulative findings in the study of HCI. What is exciting about digital libraries is the obvious attraction of such tools to most users who are exposed to them. People have always been excited by technologies that bring images, sounds, animation and text to us, but the ability to explore a huge reservoir of such information in a manner seemingly unconstrained by cost and distance is very seductive, and the growth of the internet is testimony to this fact. The convergence of the information technology, entertainment and communications industries promises a potentially seamless web of information access via the home, the office or public space.

The consequences of this attraction are important. Perhaps more than any other application in the relatively short history of information technology, digital libraries are truly a technology for the masses. The challenges then are enormous. We have here writ large the familiar problem of gaining resources for interaction issues while there are competing demands for resources to address the technical issues of copyright, bandwidth, cost, service provision etc. But there is reason for optimism here. HCI is accepted as important by anyone who sees a digital library, who watches others try to formulate a search query or try to navigate a large document space, or who examines the behavior of people in cyberspace. The potential difficulty for the field is making the transition from being seen as studying issues of importance to being seen as capable of offering the theoretical and methodological means of addressing these issues in a manner that yields answers (if provided with the resources).

The field's value, we might thus argue, is self-evident. But doubts remain. The excitement with which the Web has been received has resulted in a lot of attention being given to usability and interface design, yet much of this attention is from people unwilling to learn about the more than two decades of work this field has accumulated on interaction. Even where a Web specialist expresses interest in the history of HCI, it is not uncommon to hear comments to the effect that while those findings and data are interesting, they reflect interaction with technologies from pre-Web times, and their relevance is therefore questionable since a new technology needs new methods. I receive emails regularly asking me if any of the work on hypermedia has been replicated for Web environments, as if the findings on user navigation and comprehension of linked nodes in a standalone documents are not relevant if the documents are now part of the Internet! In extreme cases it is argued that the every concept of usability is of little value in a world where browsing and exploration render task effectiveness measures elusive. Colleagues report similar experiences. An examination of commercial and personal web sites confirms that most of the HCI findings on screen layout, color combinations and readability have been ignored or violated by many designers. In some very real senses therefore, the Web has returned the study of HCI to the 1970s where all new interface features were considered potentially useful and in need of comparison. In such a climate we must recognize the problems we face in translating our findings to practice, and learn from the mistakes we have made in not transferring HCI knowledge adequately to design practice.

Do we really know our users?

To move forward the field of HCI must package itself in a way that matters to design. This is not a new issue for HCI but the stakes have never been this high. There are many ways we can do this, and some of these we are already doing, such as increasing the opportunities for people to receive education and training in the field, and by publicizing more the importance of usability to technology acceptance, but these are slow and somewhat passive activities. I am more concerned with how HCI constructs itself as an intellectual field, and primarily, how it carves up the process of interaction, because it is in this that I believe we face a serious challenge to our wish to have greater impact.

If we consider, for example the notion of the user as she exists in HCI texts and is presented in HCI conferences we see a disjointed picture. In HCI terms, users are conceptualized at one of four levels: the physical or psychophysiological, the perceptual, the cognitive and the social. Each level of user response represents a particular granularity of analysis, with associated theoretical frameworks and methodological practices. Some have argued this distinction reflects the basic information processing characteristics of humans (see e.g., Newell, 1993), with each level necessitating different disciplinary analyses to be fully understood. Most obviously, the cognitive level is considered as covering the range of human actions in the timeframe of 0.1 to, at most, 10 seconds. Anything above the upper limit is more appropriately studied by sociological than psychological methods. Anything below the lower limit is the purview of psychophysiology.

While few HCI professionals take such limits seriously, there is a real division in HCI practice that seems to follow these divisions closely. Interface designers tend to concern themselves with task analysis and feature design, employing cognitive psychology's methods in so doing. Researchers interested in collaborative work are similarly wrapped up in frameworks that draw more on social psychology and anthropology than individual cognition. While there are many exceptions to this rule, it would seem as if Newell's carve up of human activities into time slots really does reflect some deep division in the ways we conceptualize and study users, even though it is clear that such divisions appear arbitrary and prove unhelpful to many designers who seek HCI advice and input.

Such boundaries are apparent in the way in which HCI research has progressed and is applied. If we consider the research most obviously applicable to Digital Library design for example, we find rigid partitions. Gould et al (1987) showed that reading speed differences between the media were at least partially explicable by perceptual factors, and one might think this finding would be relevant to designers and DL analysts. However, many studies of digital library design, while advocating the HCI mantra of being user-centered, adopt a social constructivist view where interpretation and negotiated meaning are considered primary, and interface variables such as image quality are captured only through the comments of participants, not through any objective assessment of the

interface. Indeed, such objective tests of interface usability are rare in DL work and have even been labeled positivistic (see e.g., Kilker and Gay, 1998). The costs of such stereotyping are high. Poor interface design will limit access and collaboration will fail to materialize. But a narrow focus on say, physical or cognitive ergonomics can blinker the HCI analyst to simple but profound issues of access that prevent interaction in the final context. In both cases, the users cause is being advocated and the designers could certainly claim to be user-centered in their approach, but the results fail to deliver the usable and useful technology we all desire because the HCI input has conceived only partial views of the user.

To be effective, HCI must overcome this problem. And it will only overcome it by developing a more holistic view of users and their interactions that does not treat each level of processing, be it cognitive, social or perceptual, as if it represented all the important user issues in full. We must be aware of the embodiment of all levels in any one user - i.e., any user exists at each of these levels in one organism, and does not herself make such distinctions in a fixed manner. While there is always a need for specialist attendance on certain levels, the majority of issues we face in DL design entail a response from HCI that is multi-leveled. The divided user perspective is a natural outcome of our history but, as we move into the 21st century, HCI should rightly be seeking its own representations. Such user analyses must find room for user learning and development, the perception of aesthetics, and the role of motivation all attributes of users and interactions that tend to be factored out of current conceptualisations. As the hottest technology in town, digital libraries provide a medium that is open to such an HCI input, and we should grasp that opportunity for holism. Digital Library users are physical beings, they are certainly behaving cognitively, and invariably are engaged in social contexts as they use the tools. Piecemeal inputs will not yield significant results, and it is the results of our efforts by which HCI will be judged.

Variables in HCI research and measurement

A second problem we face is the limited range of dependent variables that HCI considers part of its conceptual structure. We have become reasonably successful at selling

usability as an operational and dynamic construct that has certain observable indices: effectiveness, efficiency and satisfaction (although within DL circles even this is not fully understood and usability is treated not as a dynamic property of interaction but as a fixed value measured in terms of how many errors a novice user makes). Usability testers who are trained in HCI analyses routinely capture these behavioral and affective measures, aware for example that usability cannot be reduced to one simple data point such as time to complete task or rating of the interface on a Likert scale. While this is certainly an improvement over matching an interface to a set of design heuristics or asking users to state what they like about a new interface, the operational approach also induces a certain blinkering, leading us to think always in terms of performance and preference measures for specific tasks.

Where in HCI do we find people measuring variables that answer questions people outside HCI are asking of us, such as how do we attract users to our resources, and once there, what might make them stay and use it? What will bring a user back to our resource again? How do I build an interface that not only allows access but supports a richer comprehension or appreciation of the contents? These are very real design concerns, which tax us to answer some fundamental questions. What role do aesthetics play in the interactive experience? What makes material more learnable by users? Does long-term use of a digital resource affect a user's ability to read and use paper documents? Can novices learn from viewing an expert's construction of an information space? How do we determine validity or authority in shared, dispersed, and perhaps chaotic information networks? The list of interesting questions is practically endless, but what is most important, it seems to me, is that questions of this kind are central to any study of interaction but are not easily answered with the currently operationalised effectiveness, efficiency and satisfaction metrics. Within hypermedia, for example, there is an important difference between satisfaction with an interface and satisfaction with information content, and these may measure very differently in a test, but few (if any) of the standard satisfaction inventories in existence enable the discovery of this distinction. Furthermore, studies of usability tend not to take long-term measures either, yet what faith can we place in data that reflects only early learning or exposure to an interface? We know users

change over time but our experiments and usability tests are dominated by a once-off assessment of performance. While current measures are useful, they are not the only measures we could take, and they are not the only measures that would be useful to us as we seek to exploit the potential of digital technologies.

The role HCI plays in the future of DLs is partly dependent on our own ability to step outside and see ourselves as contributors. Regardless of when we start the clock, any review of HCI's 20th century contribution to technology design can only be seen as piecemeal and reactionary. I believe it could not have been otherwise, but I also believe it can be different in the future. In the final section I wish to outline a somewhat more ambitious plan for HCI and to point to at least several paths that might lead us to make an improved contribution to information systems design and implementation.

Extending HCI's remit with DLs

While I have no doubts that HCI is vital for the future exploitation of DL technologies, I am equally sure that DL's are important for the development of HCI. One reason for my optimism on this front is the sheer necessity of studying interaction in DL environments. The ubiquity of web technologies alone ensures that attention from all quarters is focused on what people do online. We can extend this to concerns with technology in the classroom, the ability of cranks and extremists to access, share, and proliferate information and viewpoints, the value of digital information to organizations, the shifting nature of mediated communication over time and distance and so forth. Everywhere we look the impact of information technology on some human endeavor is of interest. In such a climate, HCI is pushed centrestage, and even if the label 'HCI' is not attached, the issues that define our field will remain crucial. In the 21st century, people will be studying HCI regardless of whether they call themselves engineers, information architects or user experience analysts. As a field we are likely to find our approaches, theories and methods sidelined unless we engage such issues head-on.

In this section I outline several specific research questions that HCI is called upon to answer in the domain of DLs. This is clearly not an exhaustive set, but a sampling of

issues. What I attempt to show with each one is that not only could HCI make a difference here, but that each issue affords HCI a chance to extend its analysis beyond usability to enhancement, and in so doing, to extend the remit of the field to the full range of human issues surrounding information technology.

The multimedia mix and match.

At a basic level, DL interfaces require the smooth blending of multimedia. There is an inevitable HCI challenge here to understand the best means of incorporating video with text, or sound with images etc. These are classic HCI issues that we have the methods to tackle and resolve. However, embedded in this topic is a set of concerns that relate to more than layout or screen real estate (important as such topics are in determining the quality of the interactive experience).

If we accept that layout and interlinking are important issues, then it is but a short step to recognizing that information design impacts the ability of the user to comprehend the display, to recognize patterns and relationships in the material being presented, and to explore these relationships with tools for filtering, exploding, and reconfiguring the display. It is not enough that we think about efficient access to a digital resource, as HCI professionals we need to be at the forefront of envisioning means of manipulating information that enable users to extract meaning. To do that requires us to examine seriously the manner in which users learn and comprehend. And it is not enough that we draw on the existing models we have of text and graphics processing. The power of this new technology to mix multimedia in real time introduces the possibility of creating information forms that have never existed for users such as immersive environments linking sound, graphics, video and text in instantly re-configurable forms from heterogeneous sources. Might we want to consider preconfigurations for user types? Are there information overload concerns to address? Again, the questions are potentially endless, and people will turn to HCI for guidance. Our current approach to usability engineering will not be sufficiently able to guide us here, and we will need to extend the range of variables we measure to include deeper measures of active cognition.

Digital genres and the perception of information shape

As we start to examine meaning and representation of content, it becomes harder to think only in terms of individual users. We know that human communication is patterned and regular, even when it appears chaotic, and humans learn the complex rules of communication over time in both general and specific contexts. These rules are shared across users through discourse communities who evolve shared representations that support access, location, comprehension, creation, and linking of information structures. Communication is full of subtle cues that participants attune to over time, and it is only in their absence that we tend to notice the importance such cues have for successful interaction. We see the emergence of these cues in such email protocols as not shouting (using all caps) or the liberal use of quotes as contextualizers for new comments. Future HCI is likely to study how such rules and patterns give rise to communicative acts that are comprehensible, and this could unlock clues to what renders information spaces navigable by meaning as much as by physical landmarks since we must inevitably break down the division between semantics and spatiality.

The new technologies will give rise to new genres of information that only exist digitally. Indeed it seems to have already done so (Dillon and Gushrowski, 2000). We may be able to influence the form of such genres so as to optimise their value to users, and to speed up their emergence. Tools that support the reconstruction of documents tailored to the reader depending on their knowledge and interests might derive various genre forms from the same underlying information units (see e.g., Graham, 1999). Where genres reflect socio-cognitive compatibilities of communities of practice, they might provide clues to future designs, thereby serving an enhancement role in our design activities rather than evaluative one.

Learning, education and instruction

One unfortunate characteristic of 20th Century studies of HCI has been the partitioning of the field and the resultant exclusion of certain key topics from the direct concern of HCI. I consider the domain of instructional technology to be central to any informed study of

human-computer interaction, addressing as it must, issues related to information layout, sequencing, formatting, and mixing for user comprehension.

The poor showing of educational technology in much of the research on learning outcome is in part due, I believe, to the shortcomings inherent in most instructional technology designs (see Dillon and Gabbard, 1998 for a review) Not only could the field of HCI prove highly applicable to instructional technology design, but by opening HCI up to the concerns of instructional designers we will surely address core issues of concern to users and designers of much 21st Century information technology. Furthermore, DL initiatives demand this of us. It is not acceptable to those funding or implementing large DL initiatives that we partition their problems into issues that fit our existing strengths i.e., HCI professionals cannot just say "we only deal with usability ". The positive impact of DLs on user learning is a major promise of this new technology. In a culture that is increasingly seeking new ways of improving education, of providing training and lifetime learning opportunities to greater numbers of people in dispersed locations, HCI needs to be able to tackle this topic directly. I believe there is tremendous scope for HCI to study and impact designs for instruction and indeed the work has already started (see e.g., Carroll, 1998 on minimalist instruction) and the payback for us will not only be increased involvement, but a richer domain of enquiry that affords intriguing possibilities for new theoretical directions.

Intelligent IR

In an ocean of information, the ability to retrieve information on demand is undoubtedly important. While this is in many ways an old concern of information scientists, it is only a relatively new one for HCI researchers. Users will not formulate complex Boolean searches, they will be completely frustrated by hit rates that fail to discriminate among search results. The technology for filtering is improving, and associated issues of user modeling might enable us to create profiles and search engines that are more personalized and relevant. If nothing else, this is a rich vein for HCI to mine, but I suspect that the future will require even more imaginative responses from this field.

The opportunity to devise technologies that tailor their delivery to our needs, be it through know-bots, agents, filters or search engines, challenges the field to understand user requirements for information over extended periods of time. In a world of ever expanding information availability, and ever improving technical capabilities to locate, present and indeed modify information, how do we deal with the inherent desire of users to limit and filter? Will we end up with the level of specialization that first worried Bush into conceiving the memex, while at the same time enabling the free access to information as and when desired. Are we able to reconcile these two apparently contradictory demands in a single technology? Any answers to such questions, it seems to me, rely on better understanding of HCI.

Ubiquity (or we want information where we are)

The desktop is not dead, but it has spawned a family of alternative workspaces that augment it, sometimes replace it, and certainly transform it. The revolution in personal digital assistants (PDAs) had its birth pangs but seems to have found a trajectory that people enjoy. The argument is less about desktops and locations but is better understood as information access, that is, we want information where we are (less wysiwyg, and more wwiwwa). Why own a home theatre, an audio system, a home PC, a laptop and a PDA (not to mention a watch all that hardware just to tell the time! a calculator, cellular phone etc.)? How about using walls as screens, spectacles as navigation systems, and access to email from public resources as common as fast food restaurants on the non-information superhighway?

The point here of course is that HCI has always advocated technology as a means, and the greatest liberations will come from freeing us from the need for multiple workstations with their commensurate physicality. The ideal might be something like an accessible DL that could be voice activated from anywhere, using whatever physical hardware is present to project a screen, offering a personal portal to the online world, shareable on demand, and, of course, suffering no download lags or out of range messages. Or the ideal might be something very different but one thing is sure, it will be different that what we have now. Do we allow enough space for such tools to emerge from our analyses of users,

tasks and work contexts? If not, that might tell us something about what we should be asking as HCI specialists.

Conclusion

These are just some of the interesting challenges for HCI. Personalization, portability, and ubiquitous access hold promises and problems for us. Induced by developments in DL technology HCI is invited to influence information systems design on a scale that is unprecedented. The challenges to provide reliable, valid and timely inputs throughout the development life cycle of these technologies are exciting, but what we may create, as a result of this work, can only be imagined. HCI will thus need visionaries as much as evaluators, or ideally both tendencies wrapped up in the one professional. Am I falling prey to millennium optimism with this view of a new HCI, theoretically rich, engaged openly and repeatedly in the conception, design and implementation of information technologies that enhance human capabilities to learn, locate and synthesize information? Of course I am, but the reality might just be closer than it seems.

Acknowledgements

The author wishes to thank Cliff McKnight, Elaine Toms, Javed Mostafa, Jack Carroll, and a further anonymous reviewer for insightful comments on earlier drafts of this chapter.

Bibliography

Bush, V. (1945) As we may think. *Atlantic Monthly*, 176(1), 101-108.

Card, S.K., Moran, T.P. and Newell, A., (1983) *The Psychology of Human-Computer Interaction*. Hillsdale NJ: Lawrence Erlbaum Associates.

Carroll, J. (ed.) (1998) *Minimalism Beyond the Nurnberg Funnel*. Cambridge MA: MIT Press

Carroll, J (2000) *Scenario-Based Design* Boston: MIT Press.

- Davis, F.; R. Bagozzi, and P. Warshaw (1989). User Acceptance of Computer Technology: A Comparison of Two Theoretical Models, *Management Science* 35(8): 982-1003.
- Dillon, A. (1994) *Designing Usable Electronic Text: An Ergonomic Analysis of Human Information Usage*. Bristol PA: Taylor and Francis
- Dillon, (1995) Artifacts as theories: convergence through user-centered design. Proceeding of the 58th Annual ASIS Conference, Medford NJ: ASIS 208-210
- Dillon, A. (1996) Myths, misconceptions and an alternative view of information usage and the electronic medium. In J. Rouet, J. Levonen, A. Dillon and R. Spiro (eds.) *Hypertext and Cognition*. Mahwah NJ: LEA, 25-42.
- Dillon, A. and Gabbard, R. (1998) Hypermedia as an educational technology: a review of the quantitative research literature on learner comprehension, control and style. *Review of Educational Research*, 68, 3, 322-349.
- Dillon and Gushrowski, B. (2000) Is the home page the first digital genre? *Journal of the American Society for Information Science*. (2000) 51, 2, 202-205.
- Fox, E. Akscyn, R., Furuta, R. and Leggett, J. (1995) Digital Libraries. *Communications of the ACM*, 38, 4, 23-28.
- Gould, J., and Lewis, C, (1985) Designing for usability: Key principles and what designers think. *Communications of the ACM*, 28, 300-311.
- Gould, J.D., Alfaro, L., Finn, R., Haupt, B. and Minuto, A. (1987) Reading from CRT displays can be as fast as reading from paper. *Human Factors*, 29(5), 497-517.
- Graham, J. (1999) The Reader's Helper: A personalized document reading environment. *Proceedings of the ACM SCIGCHI Conference: CHI'99*, New York: ACM. 481-488.
- Kilker, J. and Gay, G. (1998) the social construction of a digital library: a case study examining implications for evaluation. *Information Technology and Libraries*, 17, 2, 60-70.
- Landauer, T. (1991) Let's get real: a position paper on the role of cognitive psychology in the design of humanly useful and usable systems. In J. Carroll (ed.) *Designing Interaction*. Cambridge: Cambridge University Press, 60-73.
- Landauer, T. (1995) *The Trouble with Computers*. Cambridge MA: MIT Press.
- Lesk, M.(1997) *Practical Digital Libraries: Books, Bytes and Bucks*. Morgan Kaufmann.

- Marchionini, G. (1999) Augmenting library services: toward the sharium. Paper presented at the International Symposium on Digital Libraries, 1999.
<http://www.ils.unc.edu/~marchsharium/ISDL.pdf>
- McKnight, C., Dillon, A. and Richardson, J. (1991) *Hypertext in Context*. Cambridge: Cambridge University Press
- Muter, P. (1996) Interface design and optimization of reading of Continuous Text. In: H. van Oostendorp and S. de Mul (eds) *Cognitive Aspects of Electronic Text Processing*. Norwood NJ: Ablex 161-180.
- Nelson, T.H. (1988) Managing immense storage. *Byte*, January, 225-238.
- Newell, A. (1993), *Unified Theories of Cognition*. Cambridge MA: Harvard University Press
- Nielsen, J. (1995) *Multimedia and Hypertext: The Internet and Beyond*. Boston: AP Professional
- Norman, K (1992) *The Psychology of Menu Selection. Designing Cognitive Control at the Human-Computer Interface* Norwood NJ: Ablex.
- Shackel, B(1959) Ergonomics for a computer. *Design*, 120, 36-39.
- Shneiderman (1980) *Software Psychology: Human Factors in Computer and Information Systems*. New Jersey: Winthrop
- Shneiderman, B. , Feldman, D., Rose, A., Ferre Grau, X. (2000) Visualizing digital library search results with categorical and hierarchical axes. Proc. Of DL00, the Fifth ACM Conference on Digital Libraries, New York, ACM, 57-68.