Spatial-Semantics: How Users Derive Shape from Information Space

Andrew Dillon

HCI Lab, School of Library & Information Science, Indiana University, Bloomington IN 47401.
E-mail: adillon@indiana.edu

User problems with large information spaces multiply in complexity when we enter the digital domain. Virtual information environments can offer 3D representations, reconfigurations, and access to large databases that may overwhelm many users' abilities to filter and represent. As a result, users frequently experience disorientation in navigating large digital spaces to locate and use information. To date, the research response has been predominantly based on the analysis of visual navigational aids that might support users' bottom-up processing of the spatial display. In the present paper, an emerging alternative is considered that places greater emphasis on the top-down application of semantic knowledge by the user gleaned from their experiences within the sociocognitive context of information production and consumption. A distinction between spatial and semantic cues is introduced, and existing empirical data are reviewed that highlight the differential reliance on spatial or semantic information as the domain expertise of the user increases. The conclusion is reached that interfaces for shaping information should be built on an increasing analysis of users' semantic processing.

Introduction

User problems with large information displays multiply in complexity as presentations move from analog to digital media. Not only can the immediately available information space expand significantly, but the elements within this space may take multiple forms (text, photographs, diagrams, sound, etc.) and incorporate animation and reconfiguration. Virtual information environments can offer 3D representations, modifiable configurations, and access to large data sets that can easily overwhelm a user's abilities to filter and represent. Compounding these issues further is the fact that the human response to information technology can shift over time, and what appears difficult to use initially might prove empowering when mastered.

One frequently discussed problem is the disorientation experienced by many users as they attempt to locate and use digital information. To date, the research response has been largely based on the analysis of visual navigational aids that

might support users' bottom-up processing of the spatial display. In the present article an alternative is considered that places greater emphasis on the top-down application of structural knowledge by the user. A distinction between spatial and semantic cues is introduced. It is argued that such a conceptualization better explains existing and emerging data. Understanding the differential reliance on spatial and semantic information could offer clues to improving user efficiency in navigation, and ultimately to improving the design of information to convey meaning to specific user communities. However, arriving at this point requires greater analysis and appreciation of the sociocognitive processes underlying the production and transmission of meaning among information users and discourse communities.

The Absence of Individual Differences in Studies of Interaction

Before addressing digital space issues directly, it is worth considering the general status of individual differences in traditional research into human–computer interaction (HCI). The paradigmatic form of HCI inquiry addresses usability by manipulating interface variables and studying user performance. Researchers might typically measure performance in terms of efficiency, effectiveness, and satisfaction to determine which features of an interface are most usable, and to ensure that any design will satisfy user requirements in a given context. This is perfectly good practice for engineering a usable artifact but it is theoretically weak as a research program because it fails to explain why interface characteristics produce the results that emerge.

Theoretical attempts to explain and predict user responses have been dominated by a handful of perspectives such as: formal model work drawing from the field of cognitive psychology (Card, Moran, & Newell, 1983), anthropological approaches to work and task performance (Suchman, 1987), and European activity theories from work study (Grief, 1991). None has produced a unifying framework, and several can be applied almost interchangeably, depending on the theorist's persuasion, to yield similarly

imprecise predictions (though the formal modeling approach has shown strong predictive power for a very restricted range of task scenarios).

Common to these theoretical approaches is the general assumption that users share many traits, abilities, dispositions, and intentions—in other words, individual differences are of little consequence beyond variance in task or computing experience [see, e.g., Nielsen's (1993) distinctions of users]. At its most extreme, the formal modeling approach treats such differences as error variance or random deviation from "ideal" performance. In this way, individual differences are closed off theoretically as insufficient to drive research.

That individual differences are too important to be ignored in the design of information technologies can be seen in the findings we have on large-scale variability in user performance. Traditional performance analyses of a range of human tasks show differences in the order of 2:1 between the best and worst performers, yet Egan (1988) reported differences between users of the order of 20:1 for common computing tasks such as programming and text editing. Similar figures have been shown for programming tasks since they were first studied (Weinberg, 1971). These data indicate that where tasks are computerized, differences between users seem to be exacerbated, not reduced. We cannot rely on training or educational interventions alone to reduce this variance or to ensure usability, so we need to consider how else to accommodate diversity through design. Ultimately, it is the technology that is more malleable than the user; thus, it is in design adaptations that our efforts should be extended.

Dillon and Watson (1996) thoroughly reviewed a century of individual differences work in psychology with a view to relating this research to systems design. They concluded that by attending to this literature, reliability and validity of user and task analyses could be significantly improved, which would support greater generalization of findings across HCI applications. They further argued that HCI research could thus build a scientific base on which to develop user classifications, moving the field beyond its current overreliance on experience or job-based criteria. Individual differences in information processing and meaning creation are central to the perspective outlined in this article.

Cognitive and Knowledge Base Differences in Users

It is worth making a distinction here between individual differences that represent deep psychological processes in the user's cognitive system (e.g., spatial ability or memory span) and those that represent differences in the knowledge base of users (e.g., task expertise). Traditionally, cognitive system differences are considered constant or at least relatively immune to change. The knowledge-base differences are more transient, and subject to alteration with training or experience. However, these processes certainly interact. Cognitive differences may directly influence knowledge-base differences. For example, low spatial ability users might always struggle with high-demand visualization tasks such that the cognitive factor places an upper bound on their performance in this context. Variance in users' knowledge bases will affect performance with a specific interface for users sharing a certain processing style (e.g., knowledge-base differences can distinguish the performance of two users with similar cognitive processing abilities; or alternatively, knowledge-base differences might minimize the effect of cognitive differences between users). Therefore, as we apply differential analyses more in the digital domain, it becomes necessary to open ourselves to the possibility that knowledge-base differences are themselves worthy of study, especially where they interact with important cognitive factors. In such cases, studying differences in user performance can offer both clues to potential cognitive factors that determine performance in context, and also provide insight into design alternatives to consider to ensure cognitive compatibility. It is primarily the knowledge-base source of differences that I believe are understudied at this time, and which I will emphasize in this article. It is the interaction of both the cognitive and knowledge-based processes that forms the heart of the spatial-semantic model outlined later.

Applying a Differential Perspective to Digital Environments

Like mainstream HCI, the experimental literature on hypermedia has tended to overlook differences between users. Yet recent reviews of the use of hypermedia by Chen and Rada (1996) and Dillon and Gabbard (1998) throw new light on the value of a differential perspective here too. Although Chen and Rada (1996) suggested there are limits to the explanatory power of general user differences in their meta-analyses of 18 experimental studies, their data showed user spatial ability to be a medium-size effect, echoing earlier findings of Sein and colleagues (1993), who reported a significant correlation between visual ability and speed of learning to use three different software applications (e-mail, modeling software, and operating systems). Dillon and Gabbard (1998) reviewed over 30 studies, and reported that learner ability and activity (as opposed to passivity in cognitive style terms) shows up in several studies of hypermedia as important determinants of learning outcome when using this technology.

Therefore, while individual differences are not widely studied, there is evidence that such differences are worthy of attention. There exists one such difference that seems central to ensuring the navigability, and hence, the usability of digital environments—the user's ability to perceive structure or shape in information space. In the remainder of this article I will outline how this ability manifests itself in use, and how we are studying it in our laboratory.

Navigation and the Perception of Shape in Information Space

A major source of cognitive overhead for users of digital information systems is the need to navigate complex information spaces. Frequently, the overhead is such that disorientation is experienced (the "lost in hyperspace" phenomenon), and users have difficulty locating required information. At best, this problem leads to increased time taken to cover material or attain acceptable comprehension rates with hypermedia (see, e.g., McKnight, Dillon, & Richardson, 1990); at worst, it can lead to the rejection of this technology on the part of users.

The literature on user navigation frequently assumes an invariant developmental sequence of landmark to route to survey knowledge, as users develop an even more detailed representation of their environment. However, there is good reason for thinking that this invariant model, albeit attractive, is oversimplistic. Theoretically, a person may develop survey knowledge independent of actually navigating a physical space through the use of external representations such as maps. Even more convincing are early data from several studies which suggest that landmark, route, and survey knowledge are each best suited to different types of tasks (Thorndyke & Hayes–Roth, 1982). Furthermore, most of the work on this form of representation tends to emerge from studies of users exploring physical or simulated physical spaces where the emergence of route and map type representations are required (e.g., Hubona, Everett, & Marsh, 1998). From the differentialist perspective, it is intriguing to note that individual differences such as spatial ability (Pellegrino, Alderton, & Shute, 1984; Stanney & Salvendy, 1995) and gender (Geary, 1995), may also systematically influence the use of different mental representations.

Plausibly, disorientation can occur through the overloading of short-term memory (STM), as users are required to remember their paths or attend to spatial markers. Overcoming STM limitations through signs and backtracking facilities is a practical strategy directly under the control of the design team. Thus, we have an extensive literature on the value and applicability of such spatial components of interfaces as menus (Norman, 1991), maps, and browsers (Simpson & McKnight, 1990), pop-up windows (Stark, 1990), direct jumping (Wright & Lickorish, 1990), etc. Such findings can work well where our goal is to support movement through an information space to a specific location, although it is cautionary that emerging data sets of user behavior in web environments reveals a greater reliance on simple "back" button navigation rather than the designed history feature of a browser (Byrne, John, Wehrle, & Crow, 1999). This feature-oriented approach to research and design derives from the classic journey metaphor of hypermedia use [see, e.g., an early manifestation of this journey metaphor in Hammond & Allinson (1987) and a current manifestation of it in Calvi (1997)].

Dillon et al. (1993) first argued that this navigation model and its attendant focus on screen features by-passes discussions of semantic space. Yet typical users are often directly interacting with the meaning of information, not consciously seeking to navigate through a space. Thus, we argued, to ensure better design for learning and communication, greater emphasis should to be placed on the semantic issues that impact use. Such arguments have been extended in recent years to the notion of information possessing shape (those spatial-semantic properties that convey coherence) that users can exploit both semantically and physically to gather meaning (Dillon & Schaap, 1996).

Dillon and Vaughan (1997) invoked shape as a construct to consider in design and evaluation of hypermedia, arguing that by so doing, one can more directly tap the variables that determine the success or failure of a hypermedia—namely, the facility with which a user can exploit the information resource to satisfice their task requirements. For users to gain from their interaction, they must do more than locate target information, at least in most real-world tasks. They must assimilate information into their own knowledge structures, an activity that extends processing beyond traversing layout to interpreting meaning.

The concept of shape assumes that an information space of any size has both spatial and semantic characteristics. That is, as well as identifying placement and layout, users directly recognize and respond to content and meaning. Routinely in our lab, users describe what they remember from an interaction in digital space or draw their interpretation of the information space's form and layout. These data clearly point to the intercoupling of spatial and semantic components of memory. For example, when asked to describe an information space after interaction, users employ terms that convey relationships and elaborations as well as purely spatial linkages such as position and sequence (see, e.g., Dillon & Vaughan, 1997). Completely separating both forms of representation is rare and somewhat artificial to users of an information space. Users easily move from one to the other, because both serve to advance their desire for task completion. Indeed, it makes best sense to think of the user's model of the information space as being constructed out of both.

It is clear from such observations that users are purposive in their dealings with information, and accessing text, graphics, sound, or images is but part of their overall requirement. Although location is necessary, it is unlikely that location itself will be sufficient in many interactive tasks. Once found, information is processed for relevance and meaning to the user. Reactions to this information cause further interactive behaviors. Furthermore, as future technologies seek to immerse users in a three-dimensional information space of multiple documents and clusters of information to enhance learning, knowledge creation and information transfer, a design target of "supporting location through navigation" will likely prove overly limited.

Futurism notwithstanding, there is a compelling reason to examine the semantic issues of use. Humans manifest a

native cognitive tendency to impose structure on information through use which is crucial to identifying appropriate information visualizations. "Shaping displays," those forms of presentation that represent spatial and semantic properties in meaningful forms, are essential if we are to move the technology from its current existence as an access mechanism to its future as a knowledge tool or an augmenter as originally envisaged by Bush, Engelbart and Nelson.

Examining the Spatial-Semantics of Information Shapes: Some Examples

How are shape perceptions formed, how persistent are they, and how can we best examine their formation and application? It is questions such as these that drive our current hypermedia research in the HCI Lab at Indiana University. In the present section, I will outline examples of this research to illuminate the points rather than report any experiment in full detail. This will provide researchers with some ideas about how these questions can be tackled. In all cases, our application domain is digital information environments that we envisage as applying at the desktop as well as immersive virtual enviroments, which we consider as a logical extension of this work.

In a series of studies exploring user's abilities to locate themselves in electronic text environments, we have shown that frequent users become attuned to regularities in form and expression of ideas that distinguish certain elements of an information space in semantic terms, and these may override spatial cues in providing users with a sense of location. For example, in a series of studies of users' ability to categorize isolated paragraphs of a scientific article presented on screen, Dillon (1991) originally showed that experts could perform this task with 80% accuracy without reading for comprehension. These text items were presented with no obvious spatial information available to the user; thus, experts must have been able to base their location on implicit signs in the discourse. It was assumed that there must exist clear perceptual cues in the discourse (keywords, formulae, etc.) that experienced readers can easily exploit to determine location because such cues might represent specific examples of categorical form.

Dillon and Schaap (1996) report the partial results of a set of follow-up studies that manipulated the presence or absence of textual cues, and examined the role of expertise with this discourse form in performing such location tasks. Results showed a significant effect for expertise, with novices manifesting greater numbers of errors than experts in terms of information location. The authors concluded that the ability to sense where one is in information space seems to be reliant in part on the rapid identification of details visually present in the information space. Experts exploit more than one source of such information to gain a clear sense of location, and thus, can categorize the information display as belonging to a certain part of a larger information space or structure. Novices or inexperienced users of an information type do not possess the necessary knowledge to interpret these cues, and therefore, must rely on explicit spatial indicators (such as headings or titles) in the visual display or on their own limited knowledge structure of the general class of information types in the world. In differential terms, experts could process semantics of the visual information space, while novices, lacking the appropriate knowledge base, had to rely solely on spatial cues.

However, the precise form of cue that experts latched on to was not simply isolated. As in the original experiment, the interface itself provided no explicit spatial cues for users to exploit, and Dillon and Shaap (1996) made a plausible case for experimentally manipulating obvious (to experts) visual cues such statistical formulae, author references, test results, descriptions of methods, etc. As such, these could be seen as representing basic semantic cues inherent in the language of this information space, which are invariant across the form. Interestingly, these authors report that the cue manipulation was not a significant effect; rather, it interacted significantly with category (i.e., superstructural form of the information space). The implication from this is that the cues manipulated were insufficient to explain the basis for the experts' superior performance.

If experts were developing a sense of location out of more than the primitive visual cues varied in these presentations, then it is clear that some nonsurface level representation in the display was affording them information on the organization of the document. To identify what these might be, the present author and a colleague (Misha Vaughan) have been examining the verbal protocols of experts as they attempt this task.

In this investigation the presence or absence of the cues originally manipulated by Dillon and Schaap (1996) were counterbalanced across categories. The subject's task at this point was to categorize each paragraph and state the reasons verbally for their choice of location.

Initial results from these types of protocols are interesting. Consider for example the following paragraph of text:

> In addition to assessing more completely the entire problem-solving process, the modified problem-solving measure has another advantage over the problem-solving measure; rather than using contrived hypothetical problems and solutions to assess problem-solving skills, it is adaptable to any idiosyncratic problem. The ability of an individual to deal with specific problems can therefore be assessed. In the present study, we asked subjects to problem-solve their response to obtaining a poor grade on an introductory psychology exam prior to obtaining such a grade. Problem-solving deficits regarding this specific stressor were expected to predict response to this stressor.

Presented like this, it possesses little spatial information, only providing an affordance to read linearly. Asked to categorize this paragraph, (which is an introduction paragraph from a scientific article, stripped of obvious cues) experts presented protocols such as these:

> Well this is interesting. My initial impression was that this is going to be discussion. And then I shifted back to meth-

ods. It's talking mostly about . . . measures. Um . . . it could be . . . . No it's introduction actually. It's because it's an overview of the method. Some studies put this in an overview of a methods section. . . . Most likely here it's the introduction, and the other thing that indicates that we're not further on, say in discussion is that it's talking about what's expected to happen. . . .

and from another expert:

Well, I am stuck. . . . Don't know where to put this one. It's either in an intro where you are setting up what you are about to tell me in more depth or you've already told me in depth and now you are presenting it to me in a discussion—so that's where I'll put it, I can see it in either. But probably . . . probably more in the discussion because it seems as if I'd know more about what this "modified problem solving measure" was . . . You would have talked about it more at length, which was earlier, So this is discussion.

So even experts disagree, particularly for introduction and discussion sections, which were previously identified as the most common mistake in such location tasks. However, without explicit cues, experts seek to base their decisions on content ("it's talking about measures") and on the anticipated form of the argument flow—"it's talking about what's expected to happen" (example 1) and "you would have talked about it more at length" (example 2). In both cases we can see that an appeal to the inherent structure of the form or genre can be made by users who are expert in this discourse type without any explicit cueing in the presentation. Such users are navigating through an information space but not in a way that is equivalent to a physical journey through space. For these users it is more a case of abstracting structure from the match between presented information and generic form that they have internalized over a lengthy socialization and training period.

In these examples, experts are exploiting knowledge shared across the community of producers and consumers. Novices are left to logically infer everything from the language without any reference to expected form or structure. Although it seems useful to design explicit spatial representations to advance the novice user's performance on such tasks, what really interests us is the extent to which the novice might gain greater semantic knowledge of the content through the representation we provide. Long term, this research is less concerned with easing access (important as that is) and more concerned with communicating knowledge and enhancing the transition from novice to more expert levels of information consumption. If, as van Dijk and Kintsch (1983) have long argued, information has a form that reflects its community's practices, we may find that designing the information space to take account of the shaping process has commensurate benefits in training new practitioners in a discipline to construct meaning.

The Spatial-Semantic Model of Information Shape Perception

Findings such as the those cited and the history of research on user navigation suggest that, ideally, the interplay of implicit and explicit spatial and semantic representations drives the formation of a working model of the information space's shape. During initial exposure to a new information space, users manifest a predictable pattern of interaction whereby they seek to find key points or landmarks within the space to which they can return. Typically, this is the top level or entry point, and users are known to repeatedly land on this point when disoriented or experiencing difficulties (Lee, Whalen, McEwen, & Latrémouille, 1984; Norman, 1991). Interestingly, this pattern emerges on paper as well as in digital information spaces (McKnight et al., 1990).

However, although most hypermedia researchers have interpreted such behavior as a motivation to pursue studies of features such as links, trails, and browsers—i.e., elements that allude to explicit spatial properties of the interface— our research places as least as much emphasis on the role of semantics. Interestingly, semantic distance has been raised as a crucial variable in understanding relevance assessment in information retrieval (see e.g., Brooks, 1995), a form of interaction that is not unlike navigation in information space where link selections must be made and followed.

Semantics also are crucial to the process through which members of a discourse community learn to shape their communications over time. Repeated interactions seem to give rise to regularities in presentation and consumption of discourse (Bazerman, 1988). Although the term "genre" has been used to cover such regularities, the definition of genre remains contested, although it points to a set of issues about information organization that would appear crucial to digital spaces (and is beginning to be addressed at last in the literature on electronic document systems; see, e.g., Reiffel, 1999). From the HCI perspective, we can conceive of communities of users learning over time to identify regularities in form that convey clues as to position in the narrative, stage of argument, likely following sequences and points of closure. Such knowledge could serve as a top-down interpretative framework that interacts with the incoming stimuli from the visual display to form an active model of information space in which the user is currently resident. This conceptualization seems in line with emerging cognitive models of reading and, in particular, the readiness effect (Gerrig & McKoon, 1998), which manifests itself in parallel processing and fast-acting resonance in memory for new input of information that has been previously processed.

The interaction of spatial and semantic processes is represented simply in Figure 1. This is a conceptual representation only, not a cognitive processing model. The information display is perceived by the user, who then creates a dynamic working model of the information space contingent on current contents and their format. Relevant spatial attributes (layout, image placement, length of text, window

FIG. 1.   The spatial-semantic model of shape perception.

size, navigation icons, etc.) combine with activated memories of just-processed information as well as semantic attributes of the information genre applied top-down (expected form, style, sequencing, meaning, etc.) to create a continuously updated and modifiable dynamic representation of the information space for this interactive task.

Seen this way, it is clear that different users will be able to apply different semantic processes to the creation and maintenance of their working model of the information space. Experience with a form or a specific space will afford such processes. Lack of experience will lead to a reliance more on spatial properties alone. Errors in performance result from the lack of certain details in the user's working model. Analyzing the type of errors made could then indicate the reliance any one user is placing on certain cues, for example, the novice's tendency to target title words in the text as indices of location in the whole document is a form of "matching bias" long known to cognitive psychologists who have studied problem solving (Johnson–Laird, 1983). Shape can, thus, be operationalized as the structural components of the working model that the user forms of an information space. Although the concept of shape this construed is most apparent in Geographical Information Systems (GIS), it is less obvious or conceptualized in textual or natural language-based environments. Such properties will always be implicit in an information space, but part of the interface design problem is to offer improved explicit representations that are compatible with the processing tendencies of humans.

Individual Differences to be Studied

Clearly, a problem for digital document designers and users is the lack of agreed genre conventions that will support the formation of shapes. One could try to transfer wholesale the existing structures of the paper world, but since the earliest days of hypertext research it has been seen that this rarely works (e.g., Wright & Lickorish, 1988). Furthermore, such copying fails to exploit the potential of digital presentation to be reconfigured according to cognitive compatibility, task dependency, or any other criterion we care to invoke. What is required, therefore, is an understanding of the properties of information that is most important in developing a sense of shape. It is important to understand the extent to which different users exploit, or find exploitable, these properties.

What becomes apparent from analyzing expert performers in the type of navigation task cited above is the complexity of the long-term process of working model formation. The surface-level semantic and visual cues manipulated over these experiments are insufficient alone to explain performance. Although the precise form of cue clearly varies across major superstructural sections of this information space, experts are exploiting deep semantics to gain a sense of order. It is this level of processing that novices clearly lack, forcing them to rely solely on surface level information that they can see or infer.

Where users of an information space possess structural knowledge of the domain, they will expect to apply deep semantic knowledge to the task in hand. Violations of conventional form or genre will detract from usability, however measured. This detraction might induce a time cost (less efficiency), an output cost (less effectiveness) or an affective cost (lower satisfaction) with the application. Similarly, total reliance on semantics to convey structure will also incur performance costs because experts also exploit spatial cues. What is interesting is the weighting of the effects—perhaps explicit spatial information can lessen the impact of structural violations at the semantic level, or perhaps semantics overwhelm explicit spatial cues. As yet, we have insufficient data to answer this, and a longitudinal study of users learning to comprehend an information space is underway to answer this.

For the novice, interaction with a new information space must necessarily be driven by spatial information. In this case, interfaces should convey explicitly the linkages, layout, and high-level organization of the space to minimize disorientation. The mainstream hypermedia literature has useful advice here, but it is insufficient for all but small and/or infrequently used spaces. All the evidence we have accumulated on the spatial-semantic issue suggests that spatial cues are coupled to semantic information as the user naturally seeks to abstract regularities in the information space. The human cognitive system continually seeks to apply existing knowledge to new information. Even with information for which there is no historical form, it is likely that genre formation starts to occur in the mind of the user almost immediately. Thus, we need to study how spatial aspects interact with semantics as the user seeks to abstract regularity from the new space. Certainly, the interaction of semantics and spatial components seems better to explain the mixed showing of spatial ability correlates with user behavior in hypermedia environments, as highlighted by Chen and Czerwinski (1997).

Beyond gross differences in the knowledge base of users, there may exist deep psychological differences based on ability or preference to deal with semantic or spatial cues or some weighted mix thereof. Such a possibility is not only suggested by the model, but is explicitly pointed at by research in education that shows differential effects for hypermedia across learner types (see Dillon & Gabbard, 1998, for a review).

Approaching the issue from the perspective of digital document designers, it is intriguing to speculate how quickly digital genres may form. If the spatial-semantic perspective is correct, it is conceivable that a genre can be formed rapidly in any situation where there exists a community of users interacting repeatedly within an information space, and this space manifests regularities that participants can perceive. Our analysis of the rapid emergence of web home pages as a genre (Dillon & Gushrowski, 2000) shows that preference for certain home page elements is positively correlated with their frequency of occurrence in pages at large. Thus, new genres are continually being formed, and the sociocognitive analysis of information shapes might offer us a cohesive means of unpacking this process.

Conclusions

The spatial-semantic model assumes all information spaces convey structural cues to the user that are differentially exploitable, depending on the user's knowledge, experience with the genre, and interactive behavior patterns. This dynamic combination of spatial and semantic information gives shape to information for the user. Although the greatest source of difference between users might be the level of semantic processing they can apply, there may be other individual differences to consider, such as the user's preference for spatial or semantic cues, which might reflect a deep cognitive style difference. This remains a concern for future research.

Designers and evaluators of digital information spaces should be informed by a clearer appreciation of the various spatial and semantic affordances manifest in any context. To this end, user analysis of the kind envisaged by Dillon and Watson (1996) could be employed to gain a more reliable and valid estimate of user differences and requirements.

The shape construct is part of a broader push for a sociocognitive analysis of interaction that seeks to reflect the interaction of multiple levels of information processing in humans, thereby blurring the rather rigid traditional boundaries between physical, perceptual, cognitive, and social perspectives that dominate current thinking. It is the present author's contention that we can only fully understand the human response to technology by adopting a multileveled analysis of interactive phenomena that vertically slices through the time-based layers of standard social science analyses. Only then can we move the field beyond designing for usability to designing for augmentation.

Acknowledgments

References

Bazerman, C. (1988). Shaping written knowledge. Madison, WI: University of Wisconsin Press.

Brooks, T. (1995). People, words and perceptions, a phenomenological investigation of textuality. Journal of the American Society for Information Science, 46(2), 103–115.

Byrne, M., John, B., Wehrle, N., & Crow, D. (1999). The tangled web we wove: A taxonomy of WWW use. Proceedings of the Annual ACM SIGCHI Conference–CHI'99 (pp. 544–551), New York: ACM Press.

Calvi, L. (1997). Navigation and disorientation: A case study. Journal of Educational Multimedia and Hypermedia, 6(3/4), 305–320.

Card, S.K., Moran, T.P., & Newell, A. (1983). The psychology of human–computer interaction. Hillsdale, NJ: Lawrence Erlbaum Associates.

Chen, C., & Czerwinski, M. (1997). Spatial ability and visual navigation: An empirical study. New Review of Hypermedia and Multimedia, 3, 67–90.

Chen, C., & Rada, R. (1996). Interacting with Hypertext: A meta-analysis of experimental studies. Human–Computer Interaction, 11(2), 125–156.

Dillon, A. (1991). Readers' models of text structures: The case of academic articles. International Journal of Man–Machine Studies, 35, 913–925.

Dillon, A., & Gabbard, R. (1998). Hypermedia as an educational technology: A review of the empirical literature on learner comprehension, control and style. Review of Educational Research, 68(3), 322–349.

Dillon, A., & Gushrowski, B. (2000) Genres and the Web: Is the personal home page the first uniquely digital genre: JASIS, 51(2), 202–205.

Dillon, A., Richardson, J., & McKnight, C. (1993) Space: the final chapter: or why physical representations are not semantic intentions. In: C. McKnight, A. Dillon & J. Richardson (Eds.), Hypertext: a psychological perspective. Chichester: Ellis Horwood, 169–192.

Dillon, A., & Schaap, D. (1996). Expertise and the perception of shape in information. Journal of the American Society for Information Science, 47, 786–788.

Dillon, A., & Vaughan, M. (1997). "It's the journey and the destination": Shape and the emergent property of genre in evaluating digital documents. New Review of Hypermedia and Multimedia, 3, 91–106.

Dillon, A., & Watson, C. (1996). User analysis HCI—the historical lessons from individual differences research. International Journal of Human–Computer Studies, 45(6), 619–638.

Egan, D. (1988). Individual differences in human–computer interaction. In: M. Helander (Ed.), Handbook of human–computer interaction (pp. 543–569), North-Holland: Elsevier.

Geary, C. (1995). Sexual selection and sex differences in spatial cognition. Learning and Individual Differences, 7, 289–301.

Gerrig, R., & McKoon, G. (1998). The readiness is all: The functionality of memory-based text processing. Discourse Processes, 26(2/3), 67–86.

Grief, I. (1991). The role of German work psychology in the design of Artifacts. In J. Carroll (Ed.), Designing interaction, psychology at the human–computer interface (pp. 203–226), Cambridge: Cambridge University Press.

Hammond, N., & Allinson, L. (1987). The travel metaphor as design principle and training aid for navigating around complex systems. In D. Diaper & R. Winder (Eds.), People and computers III (pp. 75–90), Cambridge: Cambridge University Press.

Hubona, G., Everett, S., Marsh, E., & Wauchope, K. (1998). Mental representations of spatial language. International Journal of Human–Computer Studies, 48(6), 705–728.

Johnson–Laird, P. (1983). Mental models. Cambridge: Cambridge University Press.

Lee, E., Whalen, T., McEwen, S., & Latrémouille, S. (1984). Optimizing the design of menu pages for information retrieval. Ergonomics, 27(10), 1051–1069.

McKnight, C., Dillon, A., & Richardson, J. (1990). A comparison of linear and hypertext formats in information retrieval. In R. McAleese & C.

Green (Eds.), HYPERTEXT: State of the art (pp. 10–19), Oxford: Intellect.

Nielsen, J. (1993). Usability engineering. Cambridge, MA: Academic Press.

Norman, K. (1991). The psychology of menu selection. Norwood, NJ: Ablex.

Pellegrino, J.W., Alderton, D.L., & Shute, V.J. (1984). Understanding spatial ability. Educational Psychologist, 19, 239–253.

Reiffel, E. (1999). The genre of mathematics writing and its implications for digital documents. In Proc. of the 32nd Annual Hawaii International Conference on System Sciences. Los Alamitos, CA: IEEE Computer Society (published on CD-ROM).

Sein, M., Olfmann, L., Bostrom, R., & Davis, S. (1993). Visualization ability as a predictor of user learning success. International Journal of Man–Machine Studies, 39(4), 599–620.

Simpson, A., & McKnight, C. (1990). Navigation in hypertext: Structural cues and mental maps. In R. McAleese & C. Green (Eds.), Hypertext: State of the art (pp. 73–83), Oxford: Intellect.

Stark, H. (1990). What do readers do to pop-ups, and pop-ups do to readers? In R. McAleese & C. Green (Eds.), HYPERTEXT: State of the art (pp. 2–9), Oxford: Intellect.

Stanney, K., & Salvendy, G. (1995). Information visualization: Assisting low spatial individual with information access tasks through the use of visual mediators. Ergonomics, 38, 1184–1198.

Suchman, L. (1987). Plans and situated action. Cambridge: Cambridge University Press.

Thorndyke, P.W., & Hayes–Roth, B. (1982). Differences in spatial knowledge acquired from maps and navigation. Cognitive Psychology, 14, 560–589.

van Dijk, T.A., & Kintsch, W. (1983). Strategies of discourse comprehension. London: Academic Press.

Weinberg, G. (1971). The Psychology of Computer Programming. New York: Van Nostrand Reinhold.

Wright, P., & Lickorish, A. (1988). Colour cues as location aids in lengthy texts on screen and paper. Behaviour and Information Technology, 7(1), 11–30.

Wright, P., & Lickorish, A. (1990). An empirical comparison of two navigation systems for two hypertexts. In R. McAleese & C. Green (Eds.), HYPERTEXT: State of the art (pp. 84–95), Oxford: Intellect.