Taking Artifacts Seriously

John M. Carroll, Yorktown Heights, New York

In the conventional view of theory-based design, the psychology of human-computer interaction (HCI) is seen as providing evaluations of usability or descriptive theories of the user which are then applied in the design of new computer systems and applications. This approach has not been successful in HCI. Taking HCI artifacts and situations more seriously suggests that designed systems and applications may themselves be the most useful form of 'theory' in the field. Artifacts implicitly make psychological claims (what would have to be true of users if the artifact is usable); these claims can be extracted from the artifact by a process of interpretation. This alternative resolves current methodological perplexity about the composition of the HCI field, and explains the apparent paradox of HCI development leading HCI research.

A vivid image of the recent evolution of computer technology is that of a "race" between function and usability. New technologies, new capabilities become available to users faster than user problems can be studied, understood and addressed. For example, the many user studies of word processing applications carried out over the past decade focused their attention on keyboard oriented, stand-alone systems with small and low-resolution monochrome displays. In 1981, our group at the Watson Research Center turned attention to secretaries learning to use such word processing applications. At the time, this was a novel application; computer editing was still largely the province of programmers revising code.

But now, and without a finished analysis of word processing, the frontier of usability has been pressed onward by the development and introduction of new applications and new interface technologies. Communication applications such as electronic mail and computer conference support raise usability challenges far more diverse than those raised by the extension of word processing to nonprogrammers. In the current technology, multiple users cooperatively access multiple applications via an extremely heterogeneous collection of workstation types. And even as the usability issues in these new domains are being articulated and explored, leading-edge prototypes are introducing gestural (e.g., handwriting) and speech input and interactive video output. Such new developments are occurring more rapidly, more broadly across the industry, and impacting more users all the time.

The race between function and usability creates an acute need for a *science* of human-computer interaction. If the rate of overall change in computer technology was slow, so that the appearence of new applications, new user interface designs, etc. was widely staged in time and so that new systems differed from their predecessors in only one or a few ways, then we could afford to be sanguine about "evolutionary" design: market forces would converge over time; users would eventually vote with their dollars for the better technology; good ideas would survive and poorer ones would not. However, if technology is changing on many fronts simultaneously and rapidly, then this model of evolutionary design will not work well. The interdependencies will make the situation too complex for valid attributions of cause. We are as likely to weed out good ideas as bad ones.

A science of human-computer interaction must provide means for codifying what we already know so that it can be productively applied to the problems we will have to solve in the future. It must provide a means for abstracting and generalizing what we know to help raise new empirical questions that can anticipate what we will need to know in the future. It must take us beyond casual notions of usability, beyond superficial performance analysis of systems and the people who use them to an understanding of what usability is and how to design usable systems. This is the only way we can hope to manage the race between function and usability.

1. Science and design

The relationship of basic science to design work in a technical area, for example the design of user interfaces and instruction, is one of those things that gets murkier as one examines it more closely. What is very clear and simple, however, is *why* we want to describe a close relationship between science and design. Our concrete goal is to design better solutions, better interfaces, better instruction. But we neither wish to nor expect that we can achieve this concrete goal through trial and error, through intuition or through magic: We expect that we will have to *understand* how we do what we do in design, so that we can do it deliberately and repeatedly in diverse and novel situations. Moreover, we want to be able to *externalize* our understanding of design practice to be able to teach it to others and to work with it directly to improve it.

Getting from the *why* to the *how* is the challenge. Traditional basic science seeks to develop and externalize an understanding of the world. But its primary goal is *not* to alter the world as found. On the other hand, the traditional design paradigm of craft evolution seeks to alter the world, but does not even address the problem of externalizing the inherent understanding upon which this design capability rests (Jones, 1970).

1.1 Theory-based design

There is a conventional view of the relationship of scientific research and the invention, design and development of practical artifacts. The idea is that basic science provides an understanding of nature which can be applied deductively in practical contexts. This idea has been pretty thoroughly absorbed, at least by scientists, and is familiar to anyone who has studied science. The difficulty with the idea is finding cases of invention, design and development that were driven by deduction from basic science. Practitioners know that things are not so neat; applied scientists know that invention and design *produces* theory as often as it *applies* theory.

In actual cases in which design "deductions" are offered, they are logically underdetermined. For example, Shneiderman (1980: 225) refers to George Miller's (1956) paper "The magical number seven plus or minus two" on human information processing limitations to derive the prescription that on-line training options be presented one at a time. However, there is no possible way to *deduce* this specific design guideline from the specific research and theory Miller presented on the span of absolute judgement and immediate memory. The connection is far more informal: Miller's work called attention to the (perhaps obvious) fact that humans are limited with respect to the information they can manage, but the theory he discussed was far more limited (and contentful). Shneiderman was inspired by the broader theme of limited processing capacity to suggest severely bounding the number of training options that a user ought to have to consider at a time. But this was no deduction.

The so-called systems approach to instructional design is a more extensive example of the same variety. It is remarkable to contrast the Gagne and Briggs (1979) second edition of the classic overview of the systems approach with the Gagne, Briggs and Wagner (1988) third edition. The two editions both clearly purport a deductive relationship to the psychology of learning, but they appeal to rather different views of what that psychology is: The second edition rests on Skinnerian behaviorism, while the third edition rests on the more modern information processing psychology. Amazingly both come to exactly the same instructional prescriptions. The reason this can happen is that little or no real deduction was ever involved. The systems approach to instructional design is pure methodological discipline (Jones, 1970). It has no substantive theory content and no user domain content at all. This is probably why it performs so poorly in producing instruction (Carroll, to appear).

Similar problems with theory-based design are evident in user interface design. Newell and Card (1985) outlined a "vision" for psychological science in human-computer interaction that amounts to a systems approach for theory-based design of user interfaces. They place heavy emphasis on systematic hierarchical decomposition of human behavior and experience, and on the production of simple, quantitative, time-and-error-rate descriptions. They wholly ignore the exigencies of human sense-making on the grounds that such realms of human psychology are not amenable to simple description, and hence not to design by deduction (Carroll and Campbell, 1986). This is like looking for lost car keys under a streetlight, not because the keys are anywhere nearby, but because the light seems better. Their approach has, perhaps not surprisingly, produced little impact on user interface design practice. As a remedy, they suggest pursuing "ripe application domains," and astonishingly nominate instructional systems as an exemplary domain (!).

Why doesn't the conventional view of theory-based design seem to work? The answer is partly a general fact about the relation of basic science and design work and partly a particular fact about current psychology: The science base in which design deductions must be anchored is too general and too shallow vis-a-vis specific contexts and situations in the world. In basic science, details are abstracted away; in design, they determine success or failure. Scientists want universal principles; designers need concrete examples. However, in bridging from science to design, the details cannot merely be "added back." To a great extent, the science must be redeveloped for each domain of application. Miller and Shneiderman were both concerned with processing capacity limitations, but not the *same* processing capacities or limitations. In detail, their proposals had little in common.

In the particular case of applied psychology, this mismatch of basic science and design work is aggravated by the fact that the basic work does not so much focus on abstract domains as on *odd* domains. Traditionally, academic psychology has sought to emulate physics and study abstract domains. However, subtracting the concrete meaning from domains of human experience turns out to be fundamentally unlike subtracting gravity from a physical process (in an experiment carried out in deep space). Extending the results of studies of pigeon pecking, nonsense list learning, tachistoscopic perception, etc. to the design of computer applications is hazardous at best, and often just silly. There are thousands of psychological studies of perceiving and comprehending isolated words, sentences and contrived paragraphs, but they are only of peripheral relevance to understanding real communication or, for that matter, to designing usable computer systems and instruction.

1.2 Design-based theory

The relation between basic science and invention is perplexing enough that it has often seemed reasonable to assume that they progress more or less independently. Hindle (1981) analyzed a variety of 19th century American inventions and failed to find any deductive grounding in the basic science of the time. He suggested that the conventional view of theory-based design may have developed as recently as the 1850s in the American scientific establishment essentially as a tactic for increasing the prestige of and federal support for basic research. Morrison (1974) argues that only in this century has basic science begun to exert a direct influence on practical know-how.

In the traditional paradigm of craft evolution, the knowledge designers employ and the scientific understanding they rely upon is embodied directly in their work. A 1923 book by an English wheelwright named Sturt attempted to provide explicit rationales for the manifestly successful design techniques he employed (Jones, 1970). For example, he discussed the use of outward dishing in wagon wheels: wheels are mounted so that the distance between opposing wheels is greater at the top of the wheels than at the bottom. It is notable that after a lifetime of design practice, Sturt was unable to offer a single rationale with any confidence. Rather, he produced a series of hypotheses: outward dishing in the wheels allowed the top of the wagon body to also dish outward and yet not obstruct the wheels, affording larger capacity loads; it helped to reduce the turning radius of wagons, improving maneuverability; it increased the stability of the wagon against the side-to-side lurching caused by rhythms in the horse's gait. Sturt even suggest a reason that could not possibly have been part of the original design rationale, namely, that dishing was introduced to accommodate and regularize contractions in iron tires. Iron tires were introduced long after the outward dishing of wagon wheels was established.

Technological inventions often vastly pre-date their own scientific analysis. The pulley, for example, had been used effectively for some 2,000 years before an adequate scientific analysis of its operation was developed within Newtonian mechanics. The violins of the 17th century were so finely crafted that their design was merely emulated for over 200 years. This inverted relation is also evident in current science and technology. In human-computer interaction, the development of "direct manipulation" systems (Engelbart and English, 1968; Sutherland, 1963) substantially antedates the psychological analysis of direct manipulation, and the coining of the term (Hutchins, Hollan, and Norman, 1986; Shneiderman, 1982). And, as Olson (1985) admits, current practice in designing texts far outstrips what can be grounded in the basic psychology of text comprehension.

Such inversions of theory-based design cannot be understood in the conventional view. Their resolution lies in a different view of the relation between science and design, one that takes designed artifacts and the process of invention and development that produces designed artifacts more seriously. We refer to this view as "design-based theory." In examples like the pulley, the violin, direct manipulation and text design, current understanding is embodied in the designed artifacts themselves. The artifacts provide a convenient medium for empirical exploration: unlike theoretical abstractions, they can actually be used by people in concrete circumstances. Failures can be analyzed and redesigned; successes can be emulated. In this manner, designed artifacts can serve the traditional role of theories in an applied science, codifying current understanding and guiding future efforts, but they can do this in a form more appropriate for design (Carroll and Campbell, 1989).

2. The task-artifact cycle

The perspective of designed-based theory conceives of the relation between science and design not as one-way and deductive, but as interactive and reciprocal. For example, design-based theory in human-computer interaction develops through a "task-artifact cycle" (Carroll and Campbell, 1989): People want to engage in certain tasks. In doing so, they make discoveries and incur problems; they experience insight and satisfaction, frustration and failure. Analysis of these tasks is the raw material for the invention of new tools, constrained by technological feasibility. New tools, in turn, alter the tasks for which they were designed, indeed alter the situations in which the tasks occur and even the conditions that cause people to want to engage in the tasks. This creates the need for further task analysis, and in time, for the design of further artifacts, and so on. HCI is the study of this ecology of tasks and artifacts.

An example of a task is sending a form letter to customers in Oregon. This task has an articulated structure. It involves composing, typing and revising a text, duplicating copies, putting copies into envelopes, stamping and mailing them. Analysis of the task suggests classes of artifacts that could simplify it, for example, a word processor can simplify the subtask of composing, typing and revising. However, injecting this artifact into the task situation fundamentally alters the situation itself. For example, there may be a variety of specific usability problems in adjusting to the word processor. Analyzing the task of using the word processor can suggest specific revisions in the artifact itself. And even if the word processor is unproblematic, it may restructure the constellation of subtasks: perhaps stuffing, stamping and mailing each separate envelope will now seem more tedious. Analysis of this new task situation could suggest further classes of artifacts, for example, electronic mail and network facilities. Attention can then turn to the problem of selecting only the Oregon mailing labels from a heterogeneous listing of mailing labels. This task problem may suggest yet another artifact: a database retrieval facility.

This evolutionary sequence can be observed historically and in domains far removed from human-computer interaction. Gomory (1983) analyzed the first 150 years of technology development for the steam engine to argue that the development of technology is both more complex and less predictable than the basic research from which it is seen to spring. He showed, for example, that the "revolutionary" engines of the mid-nineteenth century actually evolved through many small steps, each relying on the chance availability of a technological niche, an application in which the technology could survive and develop. For design-based theory, the availability of suitable niches for the task-artifact cycle is a requirement for the theory-building process.

The analysis and design of violins provides a further example. The modern project of understanding the violins of the 17th century required sophisticated acoustic analysis including the development of new measurement techniques (Hutchins, 1981). But to develop and assess new laws of acoustic scaling, it was necessary to also build novel instruments, the Violin Octet (Hutchins, 1967). In this project, the design work drove the basic science to new techniques and to new understanding, which was codified, tested and elaborated through the development of new artifacts. The acoustic science and the craft of violin making were so thoroughly integrated that the key members of the research team needed both sets of skills to make progress.

3. HCI artifacts as HCI theories

The artifacts which HCI produces and evaluates necessarily incorporate psychological assumptions about their usability, about their suitability for the tasks that users want to do. Chalkboard systems, for instance, have been introduced on the assumption that users already understand how to use physical chalkboards, and that the chalkboard metaphor will make such systems easier to learn and easier to use than existing systems. Such artifacts have falsifiable empirical content (Popper, 1965): chalkboard systems could turn out to have specific features that impede rather than facilitate learning and performance. By the same token, artifacts support explanations of the form, "This system feature has this consequence for usability." In these respects, *artifacts embody implicit theories of HCI.* Although explicit theory is currently scarce in HCI, artifacts are abundant, and are fulfilling many of the functions that are conventionally associated with theories.

This view of artifacts and their function in science is a novel one. Conventionally, psychological research is seen as providing *evaluations* of usability or *descriptive* theories of the user (Carroll, 1989). Neither of these conceptions, however, acknowledges the central role of user interface design in HCI research. If there is any precedent for this claim about artifacts, it would be the view that computer simulations of task performance are theory-like. Simulations are often held to embody psychological theories (Fodor, 1968; Newell and Simon, 1972). In a number of senses, simulations are the nearest neighbor to HCI artifacts. Both depend on computer technology; both embody psychological theories, but are not themselves theories; both are formal entities requiring conceptual interpretation.

There are, however, some deep differences. Simulations are used by psychologists, for specific research purposes; artifacts are used by a wide range of people to do real work. Simulations and artifacts are also interpreted in different ways. Simulations are interpreted and evaluated by criteria of *descriptive adequacy* (Chomsky, 1965): a simulation of problem-solving behavior may be judged on the basis of how closely it fits the sequence of moves in a verbal protocol, whether it predicts all and only the kinds of errors that are observed, etc. Artifacts are interpreted and evaluated by criteria of *usability.*

If artifacts are appropriate media for the expression and development of psychological theories in HCI, the question can be raised whether making the implicit theory explicit leaves the artifact with any distinctive scientific function. On a weak version of the claim, artifacts are a provisional medium for HCI, to be put aside when HCI theories catch up. On this view, we can imagine, at some point in the future, everything important about the workings and the usability properties of an artifact being extracted as an explicit theory in propositional form. Not, of course, that the theory will capture every detail of the artifact; rather, the workings of the artifact can be understood without serious distortion in terms of a central psychological theory or theories, plus some auxiliary details of "implementation."

On a strong version of the claim, artifacts are in principle irreducible to a standard scientific medium such as explicit theories. The strong version would hold, for instance, if artifacts truly cannot be understood apart from the situations in which they are used (Winograd and Flores, 1986; Suchman, 1987).

Small details of user interfaces often have a major impact on usability. Winograd and Flores (1986) and Whiteside and Wixon (1987) claim that it is impossible in principle to anticipate the effects of such details; many can only be recognized empirically.

The importance of contextual details for usability suggests that HCI may be dealing with *complex phenomena*, as in Hayek's (1967) analysis of economics. Economic phenomena are complex because they have many different kinds of determinants. More tellingly, economic phenomena are embedded in history, which Hayek regards as an unbounded, context-dependent process unfolding in time, consisting of unique events. Historical events, in effect, have an unbounded number of types. Finally, economic phenomena essentially involve human preferences, which are subjective, unpredictable, and constantly changing. Hayek concludes that economic theories must be sharply limited in predictive power. The phenomena of HCI appear to meet Hayek's criteria of complexity.

4. Interpretation

Design-based theory offers new roles and new challenges to applied scientists. There is a need for conceptual guidance in design, a need for tools that expose the theoretical claims embodied in a design. These needs demand a competent understanding of domain details. To work within the framework of design-based theory, scientists must be able to understand, indeed to generate, the examples, the designed artifacts which are the intellectual currency of the domain. And beyond this, they must develop tools for interpreting and working more effectively with artifacts. This can be done. The development of the Violin Octet is a clear success in a highly technical design domain. Wright (1978) described such a role for science in the design of texts.

Our group at the Watson Research Center is working on the problem of extracting psychological claims from HCI artifacts at a level of abstraction that is concrete enough to provide leverage within the task-artifact cycle. To do this, we consider typical user scenarios, decomposed into task analytic categories (e.g., Frese and Altman, in press; Miller, Galanter and Pribram, 1960; Norman, 1987), asking for each arena of user action what specific psychological claims the artifact is making about its users.

A simple example of an HCI artifact is the Training Wheels interface, a reduced-function training environment for a stand-alone text editor (Carroll & Carrithers, 1984; Catrambone & Carroll, 1987). The key characteristic of this interface for the purpose of articulating its psychological claims is that the training wheels design "blocks" the consequences of problematic user selections. For example, if the first-time user selects Data Merging, a message is returned that the function is not available in the training wheels interface. This simple technique has been found to facilitate initial and continuing learning (i.e., learning advanced functions like Data Merging after the training wheels are removed).

Figure 1 presents a structured interpretation for the Training Wheels interface, Specific claims are listed major arenas of user activity: mapping task goals to device-specific intentions, creating or recognizing appropriate action plans, evaluating and generalizing outcomes (Carroll and Kellogg, 1989).

- *Goals.* Error blocking embodies the claim that the mapping of real-world task goals to system goals is facilitated by filtering inappropriate goals and goal-mappings. A user who has not yet articulated an appropriate goal is blocked from prematurely engaging advanced functions like Data Merging. This implicitly guides the user toward identifying appropriate goals like typing and printing documents. The user who has already articulated an appropriate goal is blocked from mapping it inappropriately to system functions. For example, documents must be Created before they are Printed, so the selection of Print before Create is blocked. This implicitly guides the user toward correct task-device mappings.
- *Planning/Acting.* Blocking access to irrelevant functions purports to control the potential distraction of developing and pursuing erroneous plans. Because there are no false starts, action sequences are less deeply nested and relatively stereotyped across previous and subsequent attempts. This increases the chance that learners will notice what they are doing as they practice kernel scenarios (typing and printing). It becomes more likely that these action sequences will become saved in the plan repertoire.
- *Evaluation.* The reduction in the number of possible system states, in consequence of error blocking, admits of fewer possible explanations for actions and consequences in an episode, and hence constrains the user's evaluation of an interaction. Though advanced functions are blocked, users see them listed in the context of their menus. This minimal exposure to the system's full functionality seeks to support incidental learning of the scope of the full device space.

This is a simple illustration of our approach both because the Training Wheels interface is a very simple HCI artifact (embodying only the single interface technique of error blocking) and because Figure 1 considers only a few of the claims embodied in training wheels error blocking (for example, it considers no claims pertaining to user affect). Yet it suggests elementary properties of HCI artifacts that may be important: in every arena of user activity, the Training Wheels design is *psychologically overdetermined*, that is, it embodies multiple, in-

Goals	Planning/Acting	Evaluation
working opportunistically facilitates goal identification	less distraction focuses user's attention	reduced device space constrains hypotheses
working on familiar tasks facilitates goal mapping	less nested action sequences are more salient	exposure to full menus supports incidental learning
	practicing kernel scenarios integrates basic skills	

Figure 1. Psychological claims of training wheels error blocking

dependent psychological claims (Carroll and Kellogg, 1989). If this is a general property of successful HCI artifacts, it suggests their designs could never be *deduced* from simple and monolithic theortical frameworks (cf. Newell and Card, 1985).

At the Watson Research Center, we have created the User Interface Theory and Design project to explore how psychological interpretations can be produced and how they can be used in analysis and design. We are developing a core set of illustrative interpretations of artifacts such as Training Wheels and Hypercard (Wendy Kellogg). We are also trying to develop this approach in the domain of programming and software engineering, focusing on object oriented design in Smalltalk (Rachel Bellamy). We are exploring the reciprocal implications of this richer kind of theoretical description for developmental approaches to the study of skill in software domains (Robert Campbell).

5. The current perplexity

As a field of inquiry, the study of human-computer interaction is perplexing. In the midst of enormous activity and considerable technical progress, very fundamental issues remain unresolved. For example, it would seem to be axiomatic that scientific psychology has much to contribute to an understanding of HCI phenomena and to the design of HCI artifacts. However, the role of scientific psychology in HCI is in dispute.

Some theorists argue that only certain, fairly narrow conceptions of psychology can successfully be applied. Newell and Card (1985) warn that psychology might be driven out of HCI unless it can provide quantitatively predictive cognitive models. This approach focuses on relatively low-level aspects of HCI (e.g., keystroke-level methods for ideal expert performance; Card, Moran & Newell, 1983) or on simplified HCI situations (e.g., rote learning of scaled-down text editors; Polson, Kieras & Muncher, 1987). Its objective is to provide a psychological theory-base suitable for use in HCI design (Carroll & Campbell, 1986; Newell & Card, 1986).

Other theorists, in part responding to the failure of theory-based design in HCI, hold that pursuing the goal of developing cognitive science theories of HCI may *impair* progress toward usefully understanding HCI phenomena and effectively contributing to design (Whiteside & Wixon, 1987). This approach stresses the distortion and oversimplification inherent in laboratory-bound psychology and in conventional views of theory-based design. In contrast, this hermeneutic approach recommends treating situations, users and artifacts as unique instances. Understanding such instances is seen as an interactive process of unbounded interpretation: the objective is not to identify a theory-base for application to design, but to enter into a subjective process of discovery (Winograd & Flores, 1986).

Both approaches are problematic (Carroll, 1989). Deductive bridges from theory into design are dubious and vague. It is not clear that theory-based design has yet occurred on a non-trivial scale. On the other hand, bridges from hermeneutic interpretation into design decision-making are simply mystical. There is no systematic methodology, no conceptual framework, no way to objectively ground any particular experience in something more lasting or significant.

It may be simplistic to imagine deductive relations between science and design, but it would be bizarre if there were no relation at all. We believe that sufficiently rich theories of HCI artifacts and situations can directly support design activity. However, the property of psychological overdetermination discourages hope for simple, deductive bridges from theory into design. Rather, we envision a more reciprocal relation between the articulation and rearticulation of a set of psychological claims and the iterations of design.

The hermeneutic vision is correct in stressing the multiplicity of relevant interpretations of situations, users and artifacts, but too easily conflates multiplicity and infinity, settling for indeterminate subjectivity. Our view is more disciplined in assuming that there are bounds on interpretations (i.e., they are grounded in psychology and made with respect to a task analysis) and that interpretations are valuable insofar as they produce systematic and falsifiable results.

Our hope is that structured interpretations of HCI artifacts and their situations of use offer a vehicle for capturing the psychology of human computer interaction at the right level of abstraction for a design science. Taking designed artifacts more seriously as embodiments of scientific theories and results brings more of practical activity into the purview of scientific analysis. Conceiving of the task-artifact cycle as a basic structure of research activity in applied science entrains a new view of science and design. It fundamentally challenges the conventional division of labor, and directs applied science, not toward abstract or merely eccentric domains, but toward the real world.

Acknowledgement

This paper is derived from my book *The Nurnberg Funnel: Designing Minimalist instruction for practical computer skill* and from on-going collaborative work with Wendy Kellogg and Robert Campbell (as cited below).

References

Card, S.K., Moran, T.P. and Newell, A. (1983). *The psychology of human-computer interaction.* Hillsdale, NJ: Erlbaum.

Carroll, J.M. (1989). Evaluation, description and invention: Paradigms for human-computer interaction. To appear in M.C. Yovits (Ed.), *Advances in computers, Volume 28.* New York: Academic Press.

Carroll, J.M. (to appear). *The Nurnberg Funnel: Designing Minimalist instruction for practical computer skill.*

Carroll, J.M. and Campbell, R.L. (1986). Softening up Hard Science: Reply to Newell and Card. *Human-Computer Interaction, 2,* 227-249.

Carroll, J.M. and Campbell, R.L. (1989). Artifacts as psychological theories: The case of human-computer interaction. *Behaviour and Information Technology.*

Carroll, J.M. and Carrithers, C. (1984). Blocking learner errors in a training wheels system. *Human Factors, 26,* 377-389.

Carroll, J.M. and Kellogg, W.A. (1989). Artifact as theory-nexus: Hermeneutics meets theory-based design. In C.H. Lewis (Ed.) *Proceeding of CHI'89: Conference on Human Factors in Computing Systems.* New York: Association for Computing Machinery.

Catrambone, R. and Carroll, J.M. (1987). In J.M. Carroll and P.P. Tanner (Eds.) *Proceedings of CHI + GI'87: Conference on Human Factors in Computing and Graphics Interface.* New York: ACM, pages 169-173.

Chomsky, N. (1965). *Aspects of the theory of syntax.* Cambridge, MA: MIT Press.

Engelbart, D. and English, W. (1968). A research center for augmenting human intellect. *Proceedings of Fall Joint Computer Conference, 33(1).* Montvale, NJ: AFIPS Press, pages 395-410.

Fodor, J. A. (1968). *Psychological explanation.* New York: Random House.

Frese, M. and Altmann, A. (in press). The treatment of errors in learning and training. In L. Bainbridge and A. Ruiz Quintanilla (Eds.), *Developing skills with information technology.* Chichester, U.K.: Wiley.

Gagne, R.M. and Briggs, L.J. (1979). *Principles of instructional design.* New York: Holt, Rinehart and Winston.

Gagne, R.M. Briggs, L.J. and Wagner, W. (1988). *Principles of instructional design.* New York: Holt, Rinehart and Winston.

Hayek, F. A. (1967). The theory of complex phenomena. In F. A. Hayek (Ed.), *Studies in philosophy, politics, and economics.* Chicago: University of Chicago Press.

Hindle, B. (1981). *Emulation and invention.* New York: New York University Press.

Hutchins, C.M. (1967). Founding a family of fiddles. *Physics Today, 20(2),* February.

Hutchins, C.M. (1981). The acoustics of violin plates. *Scientific American, 245(4),* October, 170-186.

Hutchins, E., Hollan, J., & Norman, D.A. (1986). Direct manipulation interfaces. In D.A. Norman and S. Draper (Eds.), *User centered system design: New perspectives on human-computer interaction.* Hillsdale, N.J.: Lawrence Erlbaum Associates.

Jones, J.C. (1970). *Design methods.* New York: Wiley.

Miller, G.A. (1956). The magical number seven plus or minus two: Some limits on our capacity for processing information. *Psychological Review, 63,* 81-97.

Miller, G.A., Galanter, E., and Primbram, K. (1960). *Plans and the structure of behavior.* New York: Holt, Rinehart and Winston.

Morrison, E. (1984). *From know-how to nowhere.* Oxford: Blackwell.

Newell, A. and Card, S.K. (1985). The prospects for psychological science in human-computer interaction. *Human-Computer Interaction, 1,* 209-242.

Newell, A. and Card, S.K. (1986). Straightening out softening up: Response to Carroll and Campbell. *Human-Computer Interaction, 2,* 251-267.

Newell, A., and Simon, H. A. (1972). *Human problem solving.* Englewood Cliffs, NJ: Prentice-Hall.

Norman, D.A. (1987). Cognitive engineering - Cognitive science. In J.M. Carroll (Ed.) *Interfacing thought: Cognitive aspects of human-computer interaction.* Cambridge, MA: Bradford/MIT Press, pages 325-336.

Olson, D.R. (1985). On the design and understanding of written texts. In T.M. Duffy and R. Waller (Eds.), *Designing usable texts.* Orlando, FL: Academic Press, pages 3-15.

Polson, P., Kieras, D., and Muncher, E. (1987). Transfer of skill between inconsistent editors. Microelectronics and Computer Technology Corporation Technical Report ACA-HI-395-87, Austin, Texas.

Popper, K. (1965). *Conjectures and refutations.* New York: Harper and Row.

Shneiderman, B. (1980). *Software psychology: Human factors in computer and information systems.* Cambridge, MA: Winthrop.

Shneiderman, B. (1982). The future of interactive systems and the emergence of direct manipulation. *Behavior and Information Technology, 1*, 237-256.

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

Sutherland, I.E. (1963). Sketchpad: A man-machine graphical communication system. *Proceedings of the Spring Joint Computer Conference, 23.* AFIPS Press: Montvale, NJ, pages 329-346.

Whiteside, J. and Wixon, D. (1987). Improving human-computer interaction -- a quest for cognitive science. In J.M. Carroll (Ed.) *Interfacing thought: Cognitive aspects of human-computer interaction.* Cambridge, MA: Bradford/MIT Press, pages 337-352.

Winograd, T. and Flores, F. (1986). *Understanding computers and cognition: A new foundation for design.* Norwood, NJ: Ablex.

Wright, P. (1978). Feeding the information eaters: Suggestions for integrating pure and applied research on language comprehension. *Instructional Science, 7,* 249-312.

Author's Address:

John M. Carroll
User Interface Institute
IBM Watson Research Center
P.O. Box 704
Yorktown Heights, New York
10598 U.S.A.