Jl. of Educational Multimedia and Hypermedia (2000) **9**(2), 99-114

Bridging the Gap: Implications of Conversation for Multimedia Development

RICHARD GOLDSWORTHY

The Academic Edge, Inc.
502 South Washington Street
Bloomington, Indiana 47401, USA
rick@academicedge.net

In everyday conversation, we are constantly constructing, deconstructing, and reconstructing interfaces between ourselves and others who inhabit the world around us. This article examines this ongoing process of interface erection and teases forth implications for the design of interfaces for multimedia environments. Among the most prominent aspects of our daily conversational interface construction which hold promise for multimedia design are *interruptibility*, *chunking*, *granularity*, *divergence*, and *reconvergence.* Each of these aspects are discussed in turn and then examined as they may be applied to multimedia interface design. In the end, designers are left with a series of questions to help guide them in considering the design of their own interfaces and, hopefully, the field is left with avenues of exploration for future research.

Any point through which information of any form may flow between distinct systems constitutes an interface between those systems. Whether we are discussing the translation and transmission of machine-readable data, binary information traveling in digital format, between a computer and the network to which it is attached, or the exchange of verbal information between two human participants in a dialectic exchange, an interface between the systems is operational. Such an interface should be thought of as a

means or mode of access to information stored in either a different form from the accessing client system or in a data storage system outside of the client system. Applied to a different realm, and stripped of the technical jargon, an interface would, for example, be constituted any time two people engage in discussion. That is to say that whenever we talk with one another, we are engaged in the ongoing construction and evolution of an interface between our internal representation of information and someone else's. (This is not to discount the possibility, which cannot be explored fully here but is very interesting, of our own (re)construction of interfaces between our own varied data stores within our overall knowledge architecture.) There are, of course, many varied forms these interfaces between people may take, no two of which will necessarily be identical. In the case of human-to-human communication, the interface may be specified in terms of the specific modes by which the communication is relayed. That is to say that the communication between people may take the form of at least aural, visual, or kinesthetic signs. These signs may be asynchronous or synchronous, store-and-forward or real-time. Each instance of an interface arises through varying combinations of these potential communication characteristics. Multimedia computing environments enable, unlike any previous mode of accessing information aside from direct, face-to-face communication, the incorporation of practically all of the aforementioned characteristics.

"Conversation" has been applied to computing and, specifically, interface design in diverse ways. Walker (1990) traced conversation in computing back to what he refers to as the third generation of computing, timesharing. System administrators began to share the resources of single machines by allocating slices of processing time among users. In doing so, a form of *turn-taking* developed. "The conversational mode of interaction was the Turing test made real—the users "conversed" with the computer, just as they might with another human on a teletype-to-teletype connection" (Walker, 1990, p. 441).

Punched cards were the media of these primitive human-computer "conversations" and these cards soon were displaced by textual conversations in the form of command-line interfaces. At the same time that the medium of such communication was changing, so was the speed. Interactivity went from punching cards, submitting them for processing, and waiting for results to more immediate entering of commands and reception of results. Human conversation, in this sense of interaction, has frequently been applied as a metaphor for understanding human-machine interaction wherein the computer is treated *as if* it were a partner in a dialogue. Efforts at the textual equivalent of a conversational partner incorporate aspects of artificial intelligence, psycholinguistics, and cognitive psychology. Early efforts

to have the computer *mimic* conversation, in this metaphorical sense, include the program *Eliza*, and all efforts to create systems that would pass the Turing test. Contemporary efforts to create active agent technologies, guides, avatars, and other simulated personas represent an extension of this use of conversation as a means of designing interfaces.

This early link between text-based, turn-taking systems and conversation has led to confusion in the application of and a tendency to limit the conceptual richness of human conversation for enhancing interface design. Brennan (1990) noted that there is an unfortunate tendency to think of computer conversation as only applying in the metaphorical sense identified above: the computer as a participant in a "tit-for-tat" dialogue, usually textual. This has created differences wherein some have set conversational interfaces over and against those that are predicated upon *direct manipulation* (Schneiderman, 1987). However, when we expand our understanding of conversations beyond the limits of utterances, the *whats* of a conversation, to view, instead, conversation as situated activity, enlivened not only by the manner but also the context of the utterances, we recognize that conversation is direct manipulation, or, perhaps more accurately, direct manipulation of something may be thought of as a conversation (Brennan, 1990).

Walker (1990) argues that *world exploring* is a better metaphor for interface design than conversation. However, Walker is referring to a limited concept of conversation, the sense identified above as interacting *as if* the computer were engaged in a dialogue, rather than a broader understanding of conversation as situated activity. His view makes sense in the context of his discussion of conversation as emerging from time-sharing; however, such a statement as "when you are interacting with a computer, you are not conversing with another person. You are exploring another world" (p. 443) highlights the understanding of conversation as solely about the utterances, an understanding that needs to be enriched as Brennan and others are striving to do.

Conversation as a situated activity increases our scope to include not only what is said but, of equal and sometimes more importance, the why and the how: the meaning and the manner. All of which become bound up in what is left unsaid. Bruner (1990) describes human activity as a striving to make meaning out of experiences in daily lives, to construct and reconstruct the narrative(s) of our existence. In his discussion, Bruner points out that in any activity people are expected to act "in a manner appropriate to the setting in which they find themselves…[As] Barker put it, when people go into the post office, they behave 'post-office.'" (p. 48). Following on this observation of expectancies, Bruner cites Paul Grice's Cooperative Principle:

> Grice proposed four maxims about how conversational exchanges are and/or should be conducted—maxims of quality, quantity, and manner: our replies to one another should be brief, perspicuous, relevant, and truthful. Departures from these maxims create surplus meaning by producing what Grice calls conversational implicatures, triggers that set off searches for a 'meaning' in the exceptional, for meanings that inhere in the nature of their departure from ordinary usage (p. 48).

Grice's maxims speak most directly to what is said and left unsaid, an important aspect of conversation and of interface design, the maxims provide a checklist against which to test responses in a given situation. Good conversation meets these maxims and, importantly, is perceived as doing so by the participants in the interaction.

The emphasis on the participants' context and understanding is paramount. In a manner of Bruner's focus on activity as meaning making through narrative, Clarke and Brennan (1990) discuss mutual activity as a quest for common ground. Conversation among participants is about forming this common ground, this shared understanding. Interface design, then, can be understood as a facilitation of this creation of common ground. Laurel (1993), following upon Clarke and Brennan (1991), claims that the "notion of common ground not only provides a superior representation of the conversational process but also supports the idea that an interface is not simply the means whereby a person and a computer represent themselves to one another; rather it is a shared context for action in which both are agents" (p. 4). This development of a common ground within a shared context of action increases, then, the scope of conversation from an emphasis solely on the content of utterances and their situational appropriateness (as in Grice's maxims) to include also the mechanisms by which a common ground, a shared understanding, is formed. Little work in HCI has been done examining this latter form although there is a wealth of theory and research examining some of these mechanisms in the social sciences, particularly conversational analysis. Conversational "repair" (Schegloff, 1992; McHoull, 1990; Schegloff, Jefferson, & Sacks, 1977; Norrick, 1991) and the use of "backchannels" (Drummond & Hopper, 1993) are two particularly gravid areas of investigation (Luff, Gilbert, & Frohlich, 1990, offer a broader treatment of conversation analysis and its role in human-computer interaction).

Moving away, then, from a strict focus on the verbal turn-taking and enlarging our sphere of interest to look at conversation as situated activity, it is worthwhile to look at aspects of daily conversations that potentially have implications for interface design. The remainder of this article will examine our daily conversational "interfaces" and then apply those to broadening our understanding of human-computer interface design.

DAILY INTERFACES

If you found yourself questioning whether we really construct interfaces on a moment to moment basis as we interact with the world around us, or if you believe that such construction seems possible but wonder why we don't look at it as such, both of these reactions are artifacts of the ease with which we usually perform such erection of interfaces when we are in conversation with others. The very efficacy of our daily interfaces allows them to be constructed and deconstructed nearly invisibly. If this were not the case, we would spend far too much time considering how the interfaces in our personal communications were being constructed and not nearly enough of our processing power on the conversation itself, the information being passed over the interface. The amount of effort an individual must put forth in order to form a viable interface between his or her own representations of knowledge and those located in a data store in a different location or stored in a different manner than he or she is accustomed is often referred to as *cognitive overhead*. Building a shared understanding, Clarke and Brennan's (1991) common ground, involves such expenditure of energy on the part of all conversational participants. This initial overhead expenditure leads to, hopefully, a reduction in future efforts as the participants build their common ground. When systems, such as the human mind, evolve mechanisms for facilitating such interaction, the cognitive overhead attributable to interface construction and interpretation is decreased and information flows more readily between the two nodes, whether these are people, machines, or even, in the case of training a dog to a particular visual interface such as hand signals, animals.

Consider the case of the traveler in a foreign land: normal language structures which have served well for the formation of interfaces between participants in conversations no longer function as effectively. The interface between others and ourselves becomes highly noticeable, and we must expend cognitive effort in not only constructing the thoughts which are to be relayed but in working through the interface in order to relay those thoughts. This decreases the overall effectiveness of communication, and we consciously consider alternative forms which may be used to bridge the gap between our own particular knowledge representation and that of others—this may include the use of hands, paper, pointing to objects, or even speaking louder. The size of the representational gap which an interface must bridge, therefore, seems directly related to the differences between the participants' practiced mechanisms of interfacing; hence, the less similar an outside information representation is to those trying to access it, the more

cognitive effort would necessarily be expended on overhead, merely gaining access to the information, and the less effort can, by this very fact, be directed toward the assimilation and interpretation of the information itself.

It is granted that this second task may also be considered an act of constructing an interface. However, the secondary interface seems to be of a different nature: the first relays the information itself, ensuring some sort of equivalence between the participant's representations; the second, on the other hand, seems more of a construction of meaning within the single conversational participant potentially through the internal interfaces briefly mentioned above. While this can only be sketched for now, it is important to realize that if such a distinction between internal and external interfaces were maintained, it is essential that when trying to transmit information from one person to another, we pay close attention not only to the successful transmission of signs across the interface but also attempt to foster the construction of similar internal interfaces and information relationships, within the recipient as well. This ensures that transmitted information may take up positions in the recipient's knowledge representations or his or her meaning structures, which are approximately isomorphic to those in the sender.

CONVERSATIONAL INTERFACES

With accurate communication as a necessity in our cultural development, human society has developed a wide range of adaptive elements that comprise our toolbox of daily interface construction and alteration. By changing our depth of field to focus specifically on some of these elements of person-to-person, direct communication, essential characteristics of the tools available within a multimedia toolbox for the construction of machine-human interfaces can be explored and evaluated.

During a typical conversation between people, there is, of course, much successful and not-so-successful interface construction and adaptation. In order to maintain an adequate interface, many elements influence our daily conversations. A primary characteristic of effective interface architecture is adaptability. Rigid interfaces within face-to-face, or any form of human-to-human interface for that matter, greatly hamper communications capability. If both participants have rigid interface structures, and they are not identical, then any place where those interfaces differ will represent a disconnection, a place where information cannot be translated across the gap at all. In the extreme case, therefore, of rigid, disparate interfaces, communication comes to a halt. There is also the case of one participant being rigid and

inflexible in her or his interface while the other participant in the dialogue is adaptable. In this case, communication may occur; however, it will be entirely the result of adaptation and flexibility on the part of the second participant. Such a situation places an inordinate cognitive load on the sole flexible participant making the act of communication in this particular situation more difficult.

These extreme cases hardly ever occur in normal conversations, although we can come up with situations that might arise, such as traveling in foreign lands or being cast adrift in an inflexible bureaucratic system. More likely, participants each flex and adapt their interfaces in response to one another and within the limits of their previously learned and practiced capability to do so. In most human-to-human conversations, therefore, the amount of cognitive load is distributed across the participants and is directly related to the disparity between the participant's initial interfaces and their fund of experience working within similar situations calling for appropriate adaptations. There is nearly always some flexibility and adaptability in our human conversations as well as at least some shared cultural, or at the very least biological, components to our conversational interfaces. In the case of a human-machine interface, however, the extreme case of a rigid partner is the norm rather than the extreme exception. Without understanding the normal moment-to-moment adaptations inherent in typical human-to-human conversation, those seeking to develop interactive multimedia systems will continue to construct interfaces. These offload the majority of the cognitive processing, the interface adapting, from the machine to the human participant thereby making the task of gathering, interpreting, and incorporating the information present within the multimedia system that much more difficult for the human participant.

INTERRUPTIBILITY

If I am holding a conversation with you, for example, and I state something which is unclear, either because I have formulated the statement in a way which is either vague or poorly constructed, or because the statement is predicated upon some assumed shared fund of knowledge which is, in actuality, not available, then you, as a participant in the dialogue have a number of adaptations available to you if you want to understand my statement. The ability to interrupt the flow of our conversation and ask for clarification, through one of many means including alternative formulations and examples, is one of out most important tools for modifying our interface in order

to maintain a relatively secure link between us. This process of interruptibility is very important to maintaining communication for reasons over and above simple lack of understanding. We also interrupt when, unless constrained by situation-specific social norms, we've heard something before, when something is below our level of discourse, the old "stop me if you've heard this before," or when excited by a new direction the conversation may take, and on and on we go. Importantly, we also interrupt, or at least postpone, conversations as a result of time limitations: "gotta go, gotta go!"

Rather than jump to the implications of each of these adaptations immediately, several highly interrelated adaptations in addition to the interruptibility already mentioned will be discussed in relation to a normal human-to-human conversation. While the implications for multimedia will often be readily apparent as the adaptations are discussed in relation to typical human discourse, the application of these within a multimedia environment for effective interface construction will be made explicit immediately following the discussion of the rest of the specific adaptations.

CHUNKING AND GRANULARITY

Interruptibility is closely aligned with both the granularity of the discussion and a defining characteristic of a dialogue, time-sharing. Granularity is the level of "chunking" which is applied to information. In the case of a conversation, this often entails either the level of difficulty or the degree of conceptual abstractness at which the information is crossing the interface. I will posit that the higher the degree of granularity inherent in a given interface instantiation, the more there is an assumption by the participants of isomorphic information representation between the participants. As we saw earlier, if this assumption of isomorphism is invalid, the individual participants must engage in some cognitive processing to overcome the disparities present but unacknowledged by the interface. This processing may overload a person, particularly when the assumption of shared knowledge is large and, therefore, the gap difficult to bridge. If I, for example, assume a student understands prime factorization of numbers when I explain the concept of the least common denominator in fractions, then I am off loading additional cognitive processing to the student over and above the concept of the least common denominator. In this instance then, because of my choice of granularity (a large chunk incorporating assumed prior chunks), the student must either attempt to understand both concepts concurrently or must, in any case, employ interface adaptations such as interruptibility in order to

reduce the granularity and, therefore, the cognitive load of the transaction. On the other hand, as implied in the discussion of interruptibility above, granularity may also be too fine. That is, the information being relayed across the interface may be such that it is trivial to the person receiving the information. Students become bored, and rightly so, when the granularity of instruction is too constrained, either through too low a level of abstraction or too easy a level of difficulty. None of us wants to participate in a discussion through which nothing is new or challenging, but nor do we want to disappear in a quagmire of information far too difficult for us to incorporate—the trick to granularity is to balance between these two extremes, and this trick is accomplished, in face-to-face communication, through constant monitoring of the effectiveness of the communication through feedback and interruption, and adapting of the chunking level of the interface as needed.

None of the above would be at all possible if participants in a conversation did not, at least from time to time, interrupt one another. This intermittent interruption of the flow of a single participant's information is usually accomplished in a normally unnoticed sharing of time. One participant will pause to allow for feedback from the other. This turn-taking, or time sharing, of the channel connected to the interface, allows people participating in dialogue to adapt to the flow of conversation, to assess the accuracy of their transmission and the level of noise which has arisen, and to attempt to correct misunderstandings. Such normal interruptions in the flow may be distinguished from those necessary, and often abrupt, interruptions caused by a breakdown in the interface. Thus, by alternating between sender and receiver positions, the participants in the conversation may ensure that both of their interfaces have been constructed in such a way as the information flowing across them remains somewhat isomorphic in both content and relational structure. Such alternation, therefore, in addition to being a key component in the feedback mechanisms of an interaction, may also be considered a vital aspect of person-to-person, ongoing, interface evaluation and adaptation.

DIVERGENCE AND RECONVERGENCE

With all of this ongoing give and take, there are two other important adaptations that we as participants in interaction often utilize: we diverge and reconverge. In the roles of both sender and receiver, it is possible for participants to both follow and lead the discussion into different areas. There is no set itinerary for a discussion in a typical conversation, and an

important element in our ability to communicate effectively with one another is in our capability to follow discussions as they diverge or digress into related areas of inquiry. While often there may be an important agenda to be covered during a meeting or conversation, how that agenda is covered nearly always involves numerous digressions and reconvergences. A sender at any given moment may digress down a different path and, usually, the receiver can adapt to this change. If the interfaces constructed between the two are appropriate to the new material then they will follow without much effort expended as overhead; however, if there is, once again, a disparity in the new area, either through granularity and similar issues or through cognitive load fostered by the digression itself (some individuals may have erected interfaces which take into account their own desire to maintain a certain level of focus within the discussion; if this is the case, then such focused interfaces must be either adapted or the conversation must be reconverged onto the constructed goal of such an interface in order for it to remain effective). This aside, the interrelated abilities of participants in a human conversation to shift, gradually, or sometimes abruptly, from topic to topic within a conversation, to return to previous topics, to forestall the introduction of a topic until a more suitable moment in the conversation, and the many other similar issues of conversational flow control are some of the most fruitful adaptations for the construction of conversational interfaces available.

MULTIMEDIA INTERFACES

It should be apparent how many of these issues relate to the attempt to create efficacious multimedia interfaces which foster effective learner participation; however, since many of the adaptations often go unimplemented in software design despite, or perhaps as a result of, their the essential, but, therefore, nearly invisible or unexamined, role in normal conversation, these implications must be brought under the scope for further implementation and research. Clearly, if a learner is to interact effectively within a multimedia environment, designers and authors must do their best to implement as many of the adaptations mentioned above within their applications. Most importantly, applications must not be developed such that the majority of the processing load is off loaded from the computer to the student. Increasing the cognitive load of the learner by shifting most, if not all, of the need for flexibility and adaptability onto him or her will result in less than optimum performance on the information/task being relayed since adaptations which could have been embedded by the designer and carried out by the

computer are instead taking up precious cycles of the learner's own processor. Since a rigid interface is generally easier to code and implement, there is a great temptation for the designer to create inflexible and hierarchical (bureaucratic?) interfaces; however, if learners are to be able most effectively to construct an interface for discovering the information present in a multimedia environment, then the designer must resist this temptation since, in the end, it will hamper the student's performance.

MULTIMEDIA AND INTERRUPTIBILITY

In the case of simple interruptibility, designers of multimedia applications should take heart of the previous discussion on human-to-human interfaces and attempt to allow interruption whenever possible. Learner's should always be able to quit at any time, or at least with a very minimal time lag due to technological restraints. As with the human conversation, learner's in an interactive multimedia environment should be able to easily quit, period. No ifs, ands, or buts. There is no excuse for a program which forces learners through several pages, up several menus, or through even a page of credits at the end. One command, either a menu, an icon, or a keyboard equivalent should allow a learner to exit the program immediately—the only exception being a simple dialog box asking if the learner really wants to quit and whether or not her or his place should be saved. Of the host of programs currently available on the market, far too many non-examples of this criteria spring rapidly to mind.

We all know the feeling of either having someone walk away as we speak (either through anger or lack of attention) and most of us have walked away from a conversation (or wanted to but felt "funny" doing so); however, an application is not a human: such issues should not arise. Users should not be irritated by having to wait to leave. However, since multimedia programs are computer-based applications and, as such, can maintain an accurate pause of a "conversation" across time indefinitely, such ease of exiting should be coupled with a similar ease of return. Keeping a log of learner activities and allowing him or her to return after exiting is an essential adaptation that, parallel not only to normal human-to-human interruptibility adaptations but also digression and reconvergence, should be implemented robustly throughout an interactive module.

Aside from the issue of just plain quitting, learners should also be able to interrupt the interface at any time and to move forward or backward at their convenience, not the convenience of the programmer's code. Therefore, if a learner finds the pace or granularity of an area within the program

to be too slow or too fine (or the converses), then the learner should be able to move forward or backward as appropriate. Never, ever, have a button on screen which appears to do nothing when the user clicks on it, as often happens with forward and back buttons when background processes, such as a movie, are running—stop whatever process is holding it up and move on. If this means implementing traps on sound and QuickTime movies, or holding up screen redraws, then do so. Learners will often find themselves lost when they click a forward button a whole host of times because nothing is happening during the time an external function is occurring only to be shot five, eight, or ten screens forward when the clicks all finally register at the same time. While some of this may be somewhat unavoidable as a result of technological limitations, or skill limitations on the part of the implementer, we should do the best we can to avoid beginning any events the learner does not have the power to end, accelerate, or decelerate.

MULTIMEDIA AND GRANULARITY

Further, granularity is closely related to interruptibility. In order to adjust granularity or adapt to it, the learner must be able to interact with the interface and change the level, or be able to diverge in order to gain more information that is assumed by the level of granularity at any given moment. This can be implemented through many different methods depending upon the designer's point of view: prerequisites skills might be tested and the program implemented to adjust appropriately or to indicate a better starting place for different learners based on their previous information. This may also be entirely learner-selected with an open control structure in which participants may initially choose and then alter both the level at which the information is presented and the method through which the information is conveyed. This second would allow the user to define the interface as the user explored it and would also be the most difficult, perhaps, to implement. Another option would be to have a general structure and level intended for the average member of a given population and then to allow, through interruptibility implemented via hypertext links and self-pacing, learners to proceed through the material at an average pace which may then be increased and decreased as the learner sees fit. The links would provide additional information when needed but would not be distracting to those who do not need the finer level of granularity.

TRACKING

By implementing a tracking system, the designer could also track the frequency with which learners are going to support links (explanation and clarification links, as opposed to exploratory links which will be discussed below) and either offer the learner the option of having the links automatically incorporated within the instruction (as opposed to being in the default optional state) or expanding the links automatically without asking. As in a good conversation, the sender would adapt, therefore, to the level requisite to the receiver, but rather than being merely an imposed adaptation, it would be negotiated based on previous experience in the conversation (application). While this sort of system could be implemented readily as stated, the reverse—that is, going from too much information, too fine a level of granularity, and adapting to a user who needs less information to learn optimally would be more difficult and would have to, perhaps, be time-based (as the previous one could be as well, based on amount of time needed to answer a question or move into a new area, perhaps coupled with, when evaluation is essential, accuracy of information transmission as gathered from questions or some other form of assessment).

EXPLORATORY AND CLARIFICATORY LINKS

Mentioned only parenthetically above, exploratory links differ from those selected for clarification. While often the two types of relationships will overlap, in general links that define terms, give examples, or portray the same information in different representations, are *clarificatory*, and those which extend the information under consideration or provide related information not necessary essential to the concept under direct consideration are *exploratory* links.

This second type of link, the exploratory link, is where all the wonderful opportunity for knowledge construction related to digressions develops within multimedia applications. The ability to explore various information strands related to the present one is the great beauty of digression. When implemented, however, in order to avoid information chaos, becoming lost in the maze of data, designers of multimedia environments must provide at least one source of reference, perhaps a log of where the learner has been previously and where, in general, the learner is located within the overall information architecture present within this particular application. Preferably this information would be accessible in multiple modes, perhaps as both

a visual maps and a textual representation. Being able to move both within this metaview of the information structure of the program as well as from the "in the mix" perspective of a given area of the application environment would facilitate both learner control of the environment as well as, perhaps, the learner's cognitive structuring of the content domain(s)—as portrayed by the designer at least. These might be implemented through visual maps showing a certain number of the learner's previous locations as well as the present location, or it might be done through an ordered list of locations which the knowledge structure to which the learner has been recently (how recent in both of these cases might be left up to the learner). This ability to go into and out of digressions, to diverge and converge, should closely simulate the typical conversations described above, should be as invisible as possible, in order to reduce the cognitive overhead of moving around. Contrast the ease with which we handle books, particularly in the vital (but, alas unexamined here) markup stage and the ability to jump to well-worn dog-eared sections, with the difficulty of doing so with novels on the computer. Not the least of our troubles with books presented through the medium of the computer is the loss of the tactile quality of heft and breadth, the "knowing" in your fingertips where a specific passage is within the text not solely by intellectual recall but by the feel of the book. We have a long way to go, but we do have the model of daily conversation as a standard to which to compare our efforts.

WHAT CAN WE DO?

While there are many other applications of the conversational model to the multimedia realm, in the end the designer is responsible for the adaptability of the interface. As in any conversation, the designer should attempt to construct the program in such a way as always to keep in mind the principle that not everyone thinks about information the way they do, and it is toward this principle that all of our conversational adaptations and elements are directed. Therefore, perhaps a significant beginning for a designer would be simply to ask three basic questions when working on an interactive multimedia interface and attempt to implement the answers as best as possible within the unfortunate constraints of the authoring systems and skills available:

1. If I were in a discussion with someone about the topic I want the computer to deliver, what sorts of conversational adaptations might I or my

partner need to make in order to ensure that our personal interfaces are working effectively? In fact, would be an excellent practice to conduct a few interviews/tutoring sessions with learners in order to discover the ways in which the content may be approached (or misconstrued).
2. How can I implement as many of those types of adaptations on the server side of the conversation? In other words, how can I keep the cognitive overhead in my learner to a minimum by maximizing the flexibility and adaptability of my interface?
3. What tools and aids within the multimedia environment can I provide which would assist an "unusual" learner to adapt to my interface when the interface itself is incapable of further adaptation?

Finally, this article has necessarily treated conversation as if it were a unified construct, however, conversation is in actually a diverse set of activities, activities loosely identified as distinct genre. Narrative, the telling of stories, is one such powerful genre of conversation. Narrative too often tends to be considered as a written, non-interactive communication form, perhaps as a result of the advent of print and the distribution of stories via that medium. However, the rich traditions of oral narrative and the oral tradition in general is becoming increasingly salient as technologies (re)enrich storytelling and (re)activate the interactive nature of narrative construction. Early work by Ong (1982) on the nature of orality, its history, and its potential for society, the work of Bruner on our own narrative construction of reality—the act of making meaning—and on the development of children's story telling capabilities (Bruner, 1983, 1990, 1996), and contemporary explorations by designers (Cassell, 1998; Laurel, 1993, especially pp. 146-147; Don, 1990;), all point to the salience of narrative as a form of conversation in the design of interactivity. Additionally, systems designed for facilitating learning have been the backdrop for much of the above discussion; however, an electronic game is arguably such a system around which learning occurs and, in any event, conversation (as an activity structure if not as the activity of dialogue) and narrative certainly play a vital role in the popular success of gaming systems and, as such, benefit from exploring the principles discussed here. In the end, the various genre of conversation have much to offer computer activity design and should not be overlooked. This article represents one possible excursion through the rich intersection of these fields.

References

Brennan, S.E. (1990). Conversation as direct manipulation: An iconoclastic view. In B. Laurel (Ed.), *The art of human-computer interface design* . Reading, MA: Addison-Wesley.

Bruner, J. (1983). *Child's talk: Learning to use language*. New York: W.W. Norton.

Bruner, J. (1990). *Acts of meaning*. Cambridge, MA: Harvard University Press.

Bruner, J. (1996). *The culture of education*. Cambridge, MA: Harvard University Press.

Cassell, J. (1998). Sotrytelling as a nexus of change in the relationship between gender and technology: A feminist approach to software design. In. J. Cassell & H. Jenkins (Eds.), *From Barbie to Mortal Kombat: Gender and computer games.* Cambridge, MA: MIT Press.

Clarke, H.H., & Brennan, S.E. (1991). Grounding in communication. In L. B. Resnick, J. M. Levine, & S.D. Teasley (Eds.), *Perspectives in socially shared cognition*. Washington, DC: American Psychological Association.

Don, A. (1990). Narrative and the interface. In B. Laurel (Ed.), *The art of human-computer interface design* . Reading, MA: Addison-Wesley.

Drummond, K., & Hopper, R. (1993). Back Channels Revisited: Acknowledgment Tokens and Speakership Incipiency. *Research on Language and Social Interaction, 26*(2), 157-177.

Grice, H.P. (1989). Logic and conversation. In H.P. Grice (Ed.), Studies in the way of words. Cambridge, MA: Harvard University Press.

Laurel, B. (1993). *Computers as theatre*. Reading, MA: Addison-Wesley.

Luff, P., Gilbert, N., & Frolich, D. (Eds.). (1990). *Computers and conversation*. London: Academic Press.

McHoul, A.W. (1990). The organization of repair in classroom talk. *Language in Society, 19*(3), 349-377.

Ong, W.J. (1982). *Orality and literacy: The technologizing of the word.* London: Methuen.

Norrick, N.R. (1991). On the organization of corrective exchanges in conversation. *Journal of Pragmatics, 16*(1), 59-83.

Schegloff, E.A. (1992). Repair after next turn: The last structurally provided defense of intersubjectivity in conversation. *American Journal of Sociology, 97*(5), 1295-1345.

Schegloff, E.A., Jefferson, G., & Sacks, H. (1977). The preference for self correction in the organisation of repair in conversation. *Language, 53,* 361-382.

Schneiderman, B. (1987). *Designing the user interface: Strategies for effective human-computer interaction.* Reading, MA: Addison-Wesley.

Walker, J. (1990). Through the looking glass. In B. Laurel (Ed.), *The art of human computer interface design* . Reading, MA: Addison-Wesley.