ATARI RESEARCH MEMOS
on the subject of

INTERACTIVE FANTASY

and related topics

Brenda Laurel
Atari Sunnyvale Research Laboratory
March 1982 - November 1983

IF at the Atari Lab

The Atari Sunnyvale Research Laboratory was founded by Dr. Alan Kay in early 1982. The charter of the Systems Research Group was to do research about and develop prototypes of entertainment products and environments for the 1990's. The lab was extremely well-funded and staffed with many of the brightest and best of the Atari community, recent graduates of the Media Laboratory at MIT, and other fledgeling wizards.

This collection of papers documents my early research in interactive fantasy systems. Important contributions were made by several of the members of the Atari Systems Research Group. Eric Hulteen designed and implemented a media room environment and collaborated on much of the theoretical work. Members of the research staff, especially Michael Naimark, Susan Brennan, Scott Fisher, Chris Crawford, Alan Kay, Dr. Arthur T. Fischell, and Jim Dunion were instrumental in the development of the interactive fantasy concept. Outside the lab, other key contributors were Jeffrey Schwamberger, Clyde Grossman, Professor Donald R. Glancy of Ohio State University, Doug Lenat of Stanford University, and Ray Bradbury. My deepest thanks to all of them for their wild brainstorms, thoughtful analyses, and enduring optimism.

These papers are copyrighted by me because they were never officially published or copyrighted by Atari. They should be cited as unpublished papers. They are arranged chronologically and most have titles. Many had Atari Research Memo numbers at one time, but the numbering system has not survived the last several years. Luckily, the author has.

Readers should know that these papers represent preliminary work. The bulk of my work in this area was created after the lab's demise and is published in my Ph.D. dissertation, "Toward the Design of a Computer-Based Interactive Fantasy System", The Ohio State University, 1986. The dissertation contains a formal discussion of the theory of IF, a survey of component technologies, and a proposed system architecture. It also contains a bibliography of resource materials and other IF research.

FIRST-PERSON SIMULATIONS AS LEARNING ENVIRONMENTS

Brenda K. Laurel

March 15, 1982
Atari Sunnyvale Research Laboratory

Prepared for the NSF/NIE Conference on
Educational Computing, Berkeley,
April, 1982.

FIRST-PERSON SIMULATIONS AS LEARNING ENVIRONMENTS

INTRODUCTION

I would like to explore and expand upon the recommendation found in section 1.2 of An Agenda for Action formulated by the National Council of Teachers of Mathematics: ". . . as new technology makes it possible, problems should be presented in more natural settings or in simulations of realistic conditions." The Council has expressed the awareness that the adults of tomorrow must be equipped to solve problems and make choices in "high-technology" occupations, and that simulations are effective tools for developing the necessary problem-solving skills. I want to explore the strengths of interactive simulations as learning environments, analyze the particular strengths of first-person simulations, and discuss the powers of the form and their real-world consequences. Ultimately, I want to identify the ethical potential of simulations and examine its implications as part of the problem-solving skills that today's student will bring to bear on the real world.

INTRINSIC MOTIVATION

In his thesis, "What Makes Things Fun to Learn? A Study of Intrinsically Motivating Computer Games" (XEROX Cognitive and Instructional Sciences Series CIS-7, 1980), Thomas W. Malone supports the thory that the quality of intrinsic motivation optimizes the effectiveness of learning environments and activities. Among the key characteristics of intrinsic motivation identified by Malone are challenge, fantasy, and curiosity. In other words, an intrinsically motivating activity engages the user on those levels.

Dr. [redacted] Gattegno and others have suggested that learning is [redacted] enhanced by providing rich sensory involvement with the [redacted]ing activity -- visual, auditory, tactile, and kinesthetic. If all of the above elements are combined, they lead to the hypothesis that learning is optimized in the presence of qualities which engage the whole person in the learning activity. Stated thus, the theory is at least as old as Aristotle and has been argued successfully for centuries.

It is easy to see how an interactive computer simulation can

incorporate most of the above characteristics, including the elements of sensory stimulation, multiple levels of difficulty, and stimulants of curiosity such as novelty and complexity. Excellent models already exist in the category of games. The element of fantasy in computer simulations, however, has powers and applications beyond those usually recognized by software designers. Fantasy is seen by Singer (19xx) as a means of susstaining a high level of mental arousal or involvement, and by Arnold (1979) as a means of suspending a child's learned assessments of the nature of possibility and thus of his own limitations.

FANTASY, IMITATION, AND CHARACTER

The presence of fantasy in an educational computer simulation would seem to violate the NCTM recommendation that "problems should be presented . . . in simulations of realistic conditions." But unless the computer simulates exactly what a child of eight would do in the absence of a computer, the element of fantasy is already involved. The real question is how does fantasy work in a learning activity, and how can its powers be employed to teach the problem-solving skills needed in the "real world?"

Aristotle observed in the Poetics that humans learn by imitation. To imitate is to engage in first-person fantasy: "What if I were a flower, a cowboy, a mommy, or a starship captain?" It would seem to follow, then, that the most educationally powerful simulations are those in which the user has a first-person role; that is, the user acts as a character in the action of the simulation. But how does the simulation shape the first-person fantasy, or character, of the user?

In answering this question, the similarities between a simulation and a drama are instructive. A simulation, like a play, is an imitation of an action made up of a series of incidents. Characters move the action forward by making choices and acting upon them. In a play, the form that the character's actions will take is defined by the form of the plot -- the series of incidents that make up the whole action. The material of "stuff" of which the character's actions are made is the matrix of words and thoughts offered up to him by the world of the play. The actor, however, may not interject his own skills or judgments into the character; the imitation is "closed" by the script. A simulation is very much like a play, except that the "stuff"

available to the user for formation of his character includes not only the materials offered up by the world of the simulation, but also his own knowledge, skills, and faculties. The events that occur on the basis of his choices and actions are not fixed in a script, but flow from the interaction. In this sense, the simulation is an "open" form. The user has a double matrix of possibilities for action.

CHARACTER AND CHOICE

The ability to climb into a role or character in a simulated environment may "engage the whole person" in one imaginative stroke, presenting at once a new world of conceptual and sensory materials and a rich matrix of new possibilities. The applicability to real-world problem-solving seems clear: simulations can provide opportunities to develop the skills needed to assess new situations and to conceive of new problem-solving strategies. What may not be clear are some of the other levels of learning that are accessed by first-person simulations and the strengths and dangers inherent in them.

A great deal of work has been done to extend the levels of cognition accessed in learning activities "downward" from the computational and rational to include the intuitive and sensory. What is not acknowledged is that as we move from the abstract to the concrete, from columns of figures to contextual problems involving whole actions and whole people, we are also engaging the upper levels of analysis and choice-making in a world charged with meaning.

In the protean game of "Hammurabi," the user plays the role of king, with the task of controlling resources in order to satisfy the populace and maintain his position of power. That he must exercise particular mathematical skills and concepts is fairly obvious. In some versions of the game, he may even be encouraged to access the compter as a tool for more complex calculations, thus learning something of the uses of the computer in problem-solving activities. But what he is also learning by default in all versions of the game that I have seen is that there are acceptable levels of death by starvation which must be incorporated in the "winning" strategy.

Characters make ethical as well as expedient choices. In the imaginative leap that engages the whole person of the

user, the ethical category is opened, and what may fill it is something that the designers of the game neither perceived nor took responsibility for.

Should some other ethic have been wired into the "Hammurabi" simulation? Must teachers of mathematics now assume responsibility for identify appropriate ethics for their students? What happens when a "good" ethic is wired in?

In the game of "Star Raiders," the user must blow up as many enemy ships as he can in order to protect the galaxy. There is no "negotiation" option. There is no interaction on the ethical level. The user's choices are binary: either blow up the enemy ships and win, or don't and lose (or be perverse and blow up your own starbases). Although the ethical stance ("protect the galaxy") may be deemed acceptable, the ethical faculty has not been exercised or developed; it has simply been defined by default.

Are such scenarios really conducive to the development of problem-solving abilities that are applicable to the real world? How will they serve the genetic engineer, the nuclear physicist, or the concerned citizen of tomorrow?

FACING THE ETHICAL COMPONENT OF SIMULATIONS

The ethical component of first-person simulations is real. As an element of intrinsic motivation, it is powerful. In the skills matrix of real-world problem solving, effective use of the ethical mode is enormously important, both in the process of formulating key questions, and in the ability to visualize and evaluate consequences. The question for software designers and teachers is not whether the ethical dimension exists, but how the ethical faculty can be exercised and developed to the end of effective problem-solving in the real world.

Part of the solution is flexibility: educational simulations should allow the user to play with different ethical sets and to observe their consequences, or to formulate explicitly the ethical dimensions of the problem. Programs must be designed to recognize user strategies and demonstrate their implications. The path of possibility from choices on any level must be demonstrated in the first-person mode. Finally, the ethical component should be opened up for user interaction and discovery.

The first and most important action we can take is to become aware of the ethical dimension inherent in the form of first-person simulation. We must observe how it is defined, who defines it, and what -- whether we like it or not -- is ultimately being learned.

RESEARCH PROPOSAL: THE POETICS OF INTERACTIVE FORM

The objective of the research is to analyze the powers of certain interactive works as art. We sense intuitively that interactive works can have art-like powers--the ability to arouse emotions or to stimulate a sense of beauty, for instance. The goal is to identify those powers, and then to discover the formal and structural mechanisms through which they are realized.

The similarities between certain interactive works and dramas are striking, and they suggest that the existing body of dramatic criticism can provide a place to begin the work of analysis. We can identify a class of interactive works that are mimetic; that is, they imitate actions, things, and events that do or might occur in some real world (the worlds of Oedipus, Hamlet, and Frodo Baggins are equally "real" in the sense of dramatic probability).

We can identify a sub-category of mimetic works whose primary intentions involve entertaining, engaging, or pleasing the humans who play with them, thereby excluding works whose goals are exclusively computational, analytical, or even instructional. The resulting class may be called poetic, and it includes many games, computer-generated worlds, simulations, and a whole host of unborn genres.

Within such poetic works, we can identify elements that correspond to the primary elements of the drama:

SPECTACLE:	everything you see
MUSIC:	everything you hear
DICTION:	meaning-bearing signs and symbols
THOUGHT:	mental processes implied by what is said and done
CHARACTER:	a bundle of traits or predispositions to act in a certain way
PLOT:	the whole action being imitated

Likewise, we can investigate interactive works in terms of some traditional aesthetic criteria, such as:

MAGNITUDE:	of a size that can be taken in by the viewer, so that it may be perceived as beautiful
RHYTHM:	the pleasing arrangement of action over time
WHOLENESS:	organic unity

(next page)

The straightforward application of such critical concepts can help us assess the quality of interactive works in ways that are already different from the standards that are normally applied ("human factors" criteria, focus group results, product test criteria, etc.).

The most interesting part of the research, however, requires the formulation of new theories in areas where linear dramatic and interactive forms diverge: the realm of the user as an agent in the dramatic action. Our understanding of the drama depends upon the plot as the whole action and formal cause of the piece. By formal cause, we mean that which provides the organizing principle for all of the other elements of the piece. But what happens when the shape of the whole action is variable, depending upon the real-time choices of an intelligent agent who exists at once inside and outside the work? Is it possible to design a scheme of probability, a frame for user action, that can still provide the necessary organizing principle? Or will the intrusion of that willful agent blow the form apart? What we are seeking is a structural understanding of beauty in a made thing that accomodates the unpredictable.

The preliminary work on this question is theoretical. By playing with the dynamics of formal and material causation, we should be able to arrive at a theoretical model. That model might take the form of a mutable character, defined in terms of both dramatic probability and degrees of freedom. From there we should be able to predict the structural effects of mutable character on the piece. The next step will be to create a "living" software model that we can observe and tweak. The criteria for success will be aesthetic in nature: questions of magnitude, wholeness, and other principles of beauty can be applied. Of course, we will want to ask the obvious: "Does the piece work as art in every instance? Is it emotionally satisfying? Organic? Pleasing? Beautiful?"

For the known and potential uses of our medium, the implications are quite powerful. Perhaps we can learn how to invoke the powers of empathy, emotional arousal, and catharsis in such "open-ended" works. Can we find a real-time aesthetic that will transform our hybrid efforts into a new way of making art? Can we employ computer technology to exponentiate artistic pleasure and richness as it has informational and computational power?

I believe that we can. The form has tremendous powers that we have yet to actualize, and that actualization depends as much upon new artistic theory as upon new technology. The idea of mutable character is but one facet of our new "poetics." The introduction of that or other equally powerful concepts may allow us to systematically transform the poetics of linear drama into a comprehensive theory of interactive works as art. That is my hope.

(next page)

What form can we expect the proof of such a theory to take? A single exemplary work could be "merely" serendipitous. One thing we can do is apply our new theory, in the form of criticism, to existing works and evaluate the effectiveness of subsequent transformations. Numerous candidates exist! A more generative application of theory lies in the design of authoring environments. To some extent, every work reflects the assumptions inherent in the tools used to build it. Most computer programs reflect the logical/mathematical grammar of their authoring languages; many shortcomings in human factors are attributable to this bias. With an art-based theory of the nature of interactive works, we should be able to construct an authoring environment that emphasizes dramatic and aesthetic principles over logical/mathematical ones. Such an environment would open the medium as a creative palette for non-engineering-oriented artists. An interesting proof of our theory would be a work, employing the full power of the medium, created by Ray Bradbury, Arthur Miller, or Frank Frazetta with such authoring tools.

The user-participatory possibilities inherent in interactive form appear to obscure its status as art. When the line between artist and audience/user becomes blurred, things tend to collapse in chaos. The artist, sensing loss of control, may simply walk away from the form as a dead-end experiment. The haphazard interaction of audience members in sixties' theatrical experiments like "happenings" and some performances of the Living Theatre created just such un-beautiful, inartistic experiences; both genres have disappeared. But when the energies of artist and audience/user are woven, as thesis and antithesis, into a new artistic whole, the possibility exists for an artistic synthesis in which real formal evolution can occur. Jerzy Grotowski's Poor Theatre, for instance, casts audience members in choral roles, providing a means for dramatic integration of their participation which enhances rather than erodes the artistic power of the new theatrical form.

We brush against these questions every day in our work in human factors, game design, and even artificial intelligence, but we rarely conceptualize them as artistic questions. The tremendous body of knowledge in literary and theatrical criticism, as well as theories drawn from linguistics and structuralism, can surely inform our efforts to find "what works."

B. Laurel
12/10/82

INTERACTIVITY

Conversation with Michael Naimark, 4/6/83

IMPLICIT VS. EXPLICIT INTERACTIVITY: Naimark sees this as a major distinction, and presented the following examples:

A. A user may select one of thirty video "channels" (offerings, pieces) by pressing one of thirty buttons. The channel of his choice is then displayed on the (single) monitor.

B. A user may select one of thirty video channels by deciding which of thirty video monitors, each displaying a different channel, that he wishes to attend to.

Case (A) is "explicit;" case (B) is "implicit." A goal would seem to be to strive for systems that operate with implicit interactivity; that is, the user is never asked to step back and make an explicit choice.

SIMULTANEITY

How might implicit interactivity be expressed? Naimark noted that Andy Lippman, current head of MIT AI lab, believes that the quality of "simultaneity" is criterional of true interaction. By this he means that the user and the system must be simultaneously monitoring one another. The interaction is an ongoing collaboration, with "inputs" (lousy word) from both parties (user and system).

This is related to the idea of a system that "models" the user in some way and implements implicit (potential?) choices by noticing how the user is behaving in relation to the system--his learning or browsing styles, for instance, or his apparent skill level (in the case of a game or CAI activity). This scenario is a bit insidious.

OBSERVATIONS

I think that the implicit/explicit distinction is a useful one, but insistence on implicit interactivity as Naimark defines it might preclude some forms of overt choice-making that could be pleasurable. The distinction can be made finer in the context of interactive drama.

To have the experience of behaving like a character inside an imaginary world, the user must be allowed to make explicit choices in the PERSONA of the character. Assuming a WILLING

SUSPENSION OF DISBELIEF, then those choices are made in the first person. The operand(s) of those choices should always be within the imaginary world itself; that is, the user should not be required to have explicit interaction with the SYSTEM or to make choices at the system level.

Rather than defining choices made by the user "in character" as a subset of what we have called explicit interaction, it would be more productive to posit the existence of a third category of interaction: DRAMATIC INTERACTION. We will use the term "explicit interaction" to refer to those choices made at the system level, and not within the context of the imaginary world that the system imitates. We will use the term "implicit interaction," as we have defined it above, to refer to those cases where the user's influence on the system (and thus on the experience itself) is accomplished indirectly, without the overt expression of choice. An example of implicit interaction in the context of interactive drama would be the modification and subsequent impact of probability heuristics within the system.

(SEE color diagram, "KINDS OF INTERACTIVITY", 4/6/83)

The most important distinction among the three can be expressed as the EXPERIENCE OF PERSON-NESS. We wish to preclude "explicit interaction" because we sense that it is intrusive--that it interrupts or even destroys the user's ongoing exerience by forcing him to shift his attention to another order of problems. This is why branching interactive movies, videodiscs, or software packages feel like a kludge. Once the user is "inside" the interaction, we want him to stay there and be perfectly comfortable; in other words, we want him to be having a "first-person" experience.

In dramatic interaction, then, the user stands in first person relation to the experience. In explicit interaction, the user operates in second person relation to the system, making, as it were, imperative statements to it. In implicit interaction, the user feels himself to be in third person relation to the system, as he is not invited (or able, unless he is very clever) to attend to or control the way he is impacting the experience.

THE MIDDLES ARE BORING

Conversations with Naimark, #2
4/13/83

"The middles are boring" is a direct quote from Dr. Alan Kay, and an observation that holds true in a wide variety of contexts. In an "interactive branching" scenario, whether in computer or videodisc software, the response of the system that directly follows the user's choice/action at a "node" provides the user with a "first-person rush"--an immediate sensation of interactivity or power. Lipmann observed the phenomenon of this "rush" at the branch points in the Aspen Disc: when a user directs the system to "turn left," it is the actual turn--the transition which immediately follows the choice--that is such a pleasurable moment. The linear sequence which follows the choice node; that is, the journey through the movie-map to the next corner or node were choice is possible, is comparatively boring.

In most current interactive computer software, the same phenomenon occurs. The location of the choice "node" is, if not explicit, at least quite poorly camouflaged and readily apparent to the user. In the case of most current models of interactive fiction, for instance (the text adventure game is a good example), the user knows that once he has made his input, he must now sit back and view the results. In all but the best of systems, those results are also quite obviously predetermined, implying that the user choice has been forcibly interpreted as one of a finite set of anticipated responses.

The net effect may be expressed as a disruption in the experienced person-ness of the interaction (SEE interact.x for a discussion of kinds of interactivity). Following the choice node, the user passes quickly from a first- or second-person experience to a third-person one.

At this point it may be worthwhile to try to describe how a third-person experience feels to the user. Moving back to the grammatical model for a moment, we notice that by "third person" we mean the pronouns, he, she, it, and they. Now, in every "interactive" experience the assumption is that the user is a participant. In a third-person experience, however, there is no "I" to feel or act; there is not even a "you" that might denote the system's recognition of my presence. Either I become invisible to the system in that it is ignoring me completely, or I am suddenly dragged through the portal--through the proscenium arch, as it were--to be one of the many he-she-it entities on the other side.

It is no wonder, given the existential implications of this

predicament, that sudden shifts to the third person can cause a certain malaise. The obvious aesthetic consequences of such a disruption are but a superficial subset of the cognitive ones: the user is literally yanked into a view in which he completely disappears.

(NEXT: How videogames succeed.)

Inter Office Memo

CORPORATE DIVISION
Corporate Research

To: **Distribution**

From: Brenda Laurel

Subject: HUMAN/DOLPHIN FOUNDATION Date: 11/16/82

WILD SPECULATIONS:

Interspecies Game

Is there a video game we could imagine that a human and a dolphin could play together? Of course, HCD couldn't sell it (without an extensive marketing effort to get hardware into dolphin homes, or into all those joints where humans and dolphins hang out together). However, such a game would let us ask some real basic questions about the nature of interaction and "fun." We could also ask some esoteric ones about the nature of imitated action, semiotic thresholds, cognitive and kinesthetic response patterns, etc. etc. All this in addition to the obvious questions about interspecies communication!

Spinoff Game

Successful design/implementation of the above would give us enough data about dolphins to simulate a dolphin persona inside a video game for humans, yielding a superb learning/game product, based on a topic with tremendous popular appeal.

Thoughts

Seriously now, folks, this is just crazy enough to be great. I suggest that Sandra and I and whoever else is interested go brainstorm with the Lilly folks and come up with two proposals: one for research, and one for a consumer product. Whaddaya think?

Distribution:

Steve Arnold
Kristina Hooper
Alan Kay
Ted Kahn
Mike Liebhold
Ann Marion
Bob Stein
Sandra Williams
Steve Weyer
Jeff Sarnoff

BL:jck

The "Tool" Metaphor
Some Speculations

B. Laurel
6/20/83

The idea of computers (and/or software) as tools has been the most pervasive metaphor applied to the technology from its inception. I want to start by suggesting that "tool" as a idea does not describe the whole potential of computing power; there are kinds of computer uses and experiences that are possible which are not at all tool-like. In fact, the tool metaphor is often destrictive because it induces us to build inappropriate and dysfunctional interfaces.

Let's start by thinking about what we mean by "tools." I can think of three varieties right off the bat, and I'm sure there are more, distinguished by their functions and uses.

The most common sort of tool (at least until quite recently) is what we might call a concrete tool. Its function is to manipulate real-world objects. The desired outcome of its use is also an object or event in the real world. The function of a hammer, for instance, is to drive or extract nails; the desired outcome of its use is to construct or dismantle real-world objects. The function of a pen is to apply ink to a surface in a controlled way; the desired outcome is to create text or drawings.

A second sort of tool might be called a perceptual tool. Its function is to augment our sensory capabilities. The outcome of its use, however, may or may not be a real-world object or event. A microscope, for example, extends our visual perception. It may be used to learn more about the nature of cells ("pure" science) or to isolate a disease agent in order to affect a cure ("applied" science).

A third kind of tool can be described as a symbolic tool. It neither manipulates nor amplifies real-world phenomena, but rather operates on symbols. Like a perceptual tool, the intended outcome of its use may be "pure" or "applied." The calculator is an example.

The important thing to notice is that for each of these kinds of tools, the function is distinct from the outcome. One rarely uses a hammer for the sheer joy of randomly driving nails. If that were the intended use of a hammer, we would probably describe it as a toy rather than a tool.

The difference between a toy and a tool is just this: for a toy, there is no distinction between function (what it DOES) and outcome (what we want to accomplish by doing it). The experience of the toy's function is pleasurable in real time. The function

IS the outcome. The same thing may be said of art.

Computers and software can indeed be used as tools. A word processor, for example, has the function of creating and editing text, and the outcome of producing documents, files, code, etc. A calculation program does not significantly differ from a calculator or an abacus in terms of function and outcome. In these instances, the computer may be said to emulate or provide enhanced versions of other kinds of tools. Although such tool-like uses of computers may be functionally more complex and efficient and may produce outcomes that are much greater in magnitude than those of the tools they emulate, the basic natures of their functions and outcomes remain the same, as does the distinction between them.

Such applications are, as we know, a subset of what computers can be made to do. Some applications, like adventure games, we might describe as "toys." Others, like synthesized music, could be described as "art." If we wish to persist in defining the computer itself as a tool, then the associated functions and outcomes must encompass all such uses (more of this later).

How, then, is the tool metaphor destructive? Let's consider an example of a computer activity: the garden variety adventure game. The function of the program is to allow the user to "play" in an imaginary world. She moves through unknown spaces, usually with the goal of gathering treasures, while avoiding or overcoming obstacles and gnarly monsters. The outcome of the experience is the experience itself -- the pleasure of the activity in real time (expertise gained through repeated use serves to enhance this pleasure).

But let's look at what the interface is telling us about the functional qualities of the program. The user may press "N" to move north, for example, or "(control)lobj" to view a list of objects. If she wants to change the qualities of her magical sword, she may suspend play, view and manipulate a list of attributes, then re-enter the action of the game. In this view, the interface functions to allow the user to communicate with the program, as distinct from pursuing that "pleasurable experience" which we identified as the outcome. Thus the interface creates a tool-like distinction between function and outcome in a context where, as we have observed, they are theoretically one in the same.

This inappropriate distinction is an artifact of the tool metaphor. If we think of something as a tool, we assume its function and outcome to be distinct. The tool-like "access language" of the interface is pulled into being by the force of the metaphor. The result is a kind of estrangement, a distancing from the very experience that is the whole raison d'être of the program. This habit of mind, more than hackerliness or the "state of the art," explains the common weaknesses of the human/computer interface.

To return to an earlier question, if we try to think of the computer itself as a tool, how do we define its functions and outcomes? I suggest that its function is to represent a wide array of phenomena, including tool-like and un-tool-like activities. The outcome of its use, it seems to me, is the representation iself. This is the meta-level that I think we fail to grasp. Even when we use the tool metaphor appropriately in particular applications or activities, we invoke it inappropriately in our image of the technology itself. An example of this phenomenon is a piece of file management software being worked on by some of the people Kristina and I encountered at UCSD last week. The "interface problem" was described by one of the programmers thus: "How do we make it easy for the secretary to know what "\\'del" is good for?" Thus, even in a tool-like activity, functions are split from outcomes at the whole-system level so that the "access language" does not manage files, but rather talks to the file-management program!

The tool metaphor at the highest level dooms us to the recursive nihtmare that plagues every discussion of human/computer interface.

(Next flame: New Metaphors)

:I: Notes
Exerpts from Two Interactions

B. Laurel
7/5/83

NOTE: The text in UPPER CASE represents user input. Gestures, visual output, etc. described in the "stage directions" in parentheses.

* * * * * * * * * *

(Looking out forward view screen at long-range view of planet) MR. SPOCK?

SPOCK: Unfamiliar technology and code, captain, but the distress signal appears to have originated here (viewscreen shows magnification of equitorial region, with blinking pointer).

FALL INTO STANDARD ORBIT, MR. SULU. LIEUTENANT UHURA, KEEP SENDING THAT REPLY AND LET ME KNOW THE MINUTE YOU ESTABLISH CONTACT. MR. SPOCK, DR. McCOY, MEET ME IN THE TRANSPORTER ROOM. (Turns to Intercom) SECURITY --?

REILLY: Yes, Captain?

TWO ARMED SECURITY PERSONNEL TO BEAM DOWN WITH US -- IN FIVE MINUTES.

REILLY: Yes sir.

(Flips off Intercom, returns to first-person visual of the bridge; moves to lift; sound of doors closing behind) TRANSPORTER ROOM. (Sound of lift engaging)

IF THAT DISTRESS CALL IS GENUINE, WE ARE GOING TO HAVE TO INTERVENE IN THEIR CIVIL WAR TO AVOID THE COMPLETE DESTRUCTION OF ONE OF THE SENTIENT SPECIES ON THE PLANET.

"I don't have enough information to consider violating the prime directive -- better not pursue this line of thought until I get down there and see for myself."

RIGHT. (Lift doors open; visual scan of corridor; walks to transporter room. SPOCK, McCOY, and SECURITY GUARDS are already in place to beam down. Walks to transporter platform). ENERGIZE. . .

* * * * * * * * * * * *

BUILD A WORLD, PLEASE . . .

"Right. I'm listening."

USE WHAT YOU KNOW SO FAR ABOUT MY FAMILY AND GIVE ME WHAT YOU CAN PUT TOGETHER ON 1966. IN THE KITCHEN, EVENING, MOM AND DAD AND ME SITTING AROUND THE TABLE.

"Right. I have personality information on your parents and I know something about what that kitchen looked like. How's this?"

GOOD. MAKE THE TABLE A LITTLE BIGGER, AND MAPLE, NOT WALNUT . . . THE WINDOWS ARE RIGHT BUT THE CURTAINS SHOULD BE ORANGE, AND CLOSED . . . SHE ALWAYS KEEPS THEM CLOSED.

"I'll remember that." (Screen image is adjusted)

HERE'S THE SITUATION. JIM AND I HAVE BEEN OUT IN MY CAR LAST NIGHT. THIS MORNING MOM FOUND HIS WALLET AND COMB IN THE BACK SEAT . . .

"Excuse me, can you tell me something about Jim?"

OH YES . . . JIM MARTZ . . . YOU SHOULD BE ABLE TO FIND HIM IN THE DIARY . . . HE WAS MY STEADY BOYFRIEND SINCE '64 . . .

"Right. Go on . . . "

SO MOTHER FINDS THIS STUFF AND DECIDES THAT WE ARE SCREWING. SHE DECIDES TO FORCE US TO GET MARRIED. HERE'S THE DIALOGUE AS I REMEMBER IT . . .

(Reports a 5-minute dialog)

YOU GOT [illegible]?

Yes. What do you want to do?

OKAY, PLAY ME WHAT YOU HAVE AND I WANT TO JUMP IN ABOUT THE TIME SHE SAYS "WE WILL GIVE YOU YOUR BEDROOM FURNITURE AND TWO THOUSAND DOLLARS . . ."

Okay. Ready?

(Dialog begins, goes on for a few minutes)

MOTHER: . . . and we will give you your bedroom furniture and two

thousand dollars . . .

WAIT A MINUTE, MOM. I DON'T THINK IT HAS TO WORK OUT THIS WAY. WHAT WOULD IT TAKE FOR YOU TO FEEL BETTER ABOUT THIS . . . ?

(At this point the dialog begins to diverge from the original scenario . . .)

* * * * * * * * * *

Notes on Interestingness in Interactive Stories: Some Playwriting Principles

B. Laurel
Created 8/11/83
Last Revised 8/12/83

Overview

The idea of attaching an intelligent playwriting system to an extant story generator is attractive for several reasons. The problems of representing characters and providing the means for them to make choices that are manifest in action are largely solved, at least conceptually, in a system like Meehan's Tale-Spin. Further, the character and thought levels of Tale-Spin are already equipped to listen to the advice of the "Storyteller" portion of the program, which currently functions, in one mode, to apply story-based criteria to clip redundant, unimportant, or boring detail from the output story and, in another mode, to synthesize a story to exemplify a fable or moral. The next logical step is to replace (or transform utterly) the Storyteller with a program which aims to alter the output of the system in two directions: to enable a user-character interface, and to employ dramatic techniques to augment "interestingness."

To implement some important dramatic techniques, it will also be necessary to expand the scope of human-authored materials upon which the system may operate, specifically to provide materials which are generative of accidental incidents and of dramatic reversals. Although we may wish later to develop ways to automate the authoring of materials (technially, the specification and delimiting of POTENTIAL), assuming that such authoring will be "hand-done" for the moment will allow us to focus upon the design and implementation of a playwriting expert as our first order of business.

This paper will hopefully ask more questions than it answers.

Theory of the User-Character Interface

Tale-Spin deals with the generation and telling of stories and, as such, uses a narrative form: the stories are episodic in construction, and are meant to be read by a non-interacting user. Once the second-person transactions with the system that are necessary to initiate action have been completed, the user's experience of Tale-Spin is essentially in the third-person mode. It is important to recognize also that the user need only establish the initial conditions of the story; the system

operates independently to synthesize the story from these inputs.

The interactive system which we envision is different in several crucial aspects. We wish to enable the user to move about in the "story" as one of its characters. First and most important, this would transform the user's experience into a first-person mode. The operation of the story generator itself would also be fundamentally changed. Currently, Tale-Spin allows the user to describe characters with traits and goals. The system then creates those characters and those features of an environment which are necessary to them, AND NOTHING ELSE. The story is generated from these materials, using the planning and action structures that are part of the "generator" itself.

Allowing the user to move about in the story as a character means that the user-character's choices and actions must be taken into account at every moment, as they are continually affecting the states, goals, and actions of other characters and the state of the "world" at large. What must the system do to accomodate this stream of inputs?

First, it must provide a means for the other (generated) characters to form models of the user-character. Each character's model will be based on the information about the user-character to which he has been exposed, as well as upon his own traits and predispositions (literally, what he is inclined and equipped to notice). Second, the system itself (specifically, the "playwriting" program, of which more later) must maintain and update a model of the user-character in order to know what is ultimately "true." This is a harder problem than it seems, for without sophisticated user constraints, the system is theoretically no better equipped than any of its characters to know the actual states and goals of the user-character.

Further complications arise from the fact that the user-character does not have the same organic relationship to the potential embodied in the story world as one of the characters that has been created with it. By their very nature, system characters are incapable of introducing new potential into the story world -- they can only formulate that potential into possibility in the pursuit of their goals. But what is to keep the user-character from introducing new potential (that is, materials which are unknown to the system in its current configuration) and thereby blowing the system up? The essential user-character problem is one of information: how can the story world (system) and the user-character exchange necessary and reliable information without violating the first-person imperative (that is, without resorting to second-person techniques like "question-and-answer" and "error messages")?

Three distinct solution strategies can be used to address the

various aspects of the problem. The problem of communicating information to the user-character about the potential of the story world in a first-person mode can be re-cast in cognitive terms: how does the user know what is possible, or "where the edges are?" The concept of EXPOSITION, borrowed from the drama, suggests a solution that operates in two ways. In plays, "exposition" is the presentation of materials and information which define the limits of the world of the play (context). The total of expository materials is made available to the audience, so in a way, the audience has the same kind of global access to context information as the playwriting program would have in our system. Exposition is also relevant to the specific characters in the play. Many of the interesting complications in a drama result from the differences in content (completeness and reliability) as well as order of presentation of expository information to various characters. This character-specific exposition corresponds to the distinctions in knowledge about the story world held by the various characters in our system.

In a play, exposition is an ongoing process, operating through and around the dramatic action, with a higher concentration of exposition occurring at the beginning of the play. Exposition would occur similarly in our interactive system, with the highest concentration at the beginning of the interaction. We might propose an "expository interval" before the dramatic action begins, during which the user would simply observe the "world in progress," gaining information about the context and characters. This strategy assumes that the entrance of the user-character into the story world functions as a sort of inciting incident, or formal beginning, of the dramatic action. Later, expository material could be doled out to the user-character as well as to the other system characters, according to access opportunities (being in the right place at the right time), goal to know (as Meehan's characters can question one another, for instance), and the potential for each morsel of information to create interesting dramatic situations (determined by the playwriting program).

A second technique for creating cognitive boundaries is to control which assertions on the part of the user-character are acknowledged by the system and become manifest in the action. If the user-character attempts to introduce or operate on potential that is beyond the boundaries of the story world, that assertion can simply fail to produce results. If user-character Joe Bear attempts to solve a problem by asserting that he has magical powers that will enable him to find honey, for instance, the action subsequent to that assertion can simply prove him wrong -- at the same time supplying the information that such powers are not part of the potential of the story world. This technique, while obvious and simple, is less intrusive and destructive of first-personness than a second-person question or error message.

(It also raises the whole issue of the relationship between degrees of freedom and the potential for pleasurable experience, which is beyond the scope of this paper but will be treated later).

(3) A third facet of the user-character problem is how the system can construct a model of the user-character for its playwriting operations which has a maximal degree of reliability. Of course, the two techniques for informing the user about the potential of the story world will help to prevent accidental errors on the part of the user in communicating his own goals, traits, and strategies, but will do nothing to preclude intentional misrepresentation. Little attention can or should be paid to intentional misrepresentation as vandalism, since human perversity will always find a way to assert itself anyway, but misrepresentation for dramatic ends (as part of a strategy for achieving a character's goal, for instance) should be accomodated and utilized as elegantly as possible.

The technique which the playwriting program might use is one of successive approximation, making and updating a "best guess" version of the user-character model whenever the user-character acts (action in this sense includes both sayings and doings). That best-guess model can then be used to orchestrate the lines of dramatic probability in the story world, and less reliable versions may also be used to create and maintain alternative probability structures (which will in turn enhance both the complexity of the plot and the richness of dramatic alternatives available to the playwriting program). When certain elements of the user-character model are shown to be erroneous, the relevant lines of probability can be re-directed on the fly, and the alternate lines can be examined for their usefulness in light of the new information. How lines of probability are configured will be discussed below.

*** INCOMPLETE DRAFT ***

(continuation of Interestingness paper)

Last revised: 8/12/83

Dramatic Probability: Introduction

Dramatic probability may be defined as the causal relations among incidents in a play. In review, we know that a play is an imitation of an action and that it has a beginning, a middle, and an end. While a play is constructed of many actions (and their agents), it is ultimately an imitation (or representation) of a single, higher-order action which unifies the whole, and it is the progression of this superordinate action which provides the beginning, middle, and end. Dramatic probability is the means by which the action moves forward.

We can describe all the potential for dramatic action which exists at the beginning of a play as a world of POSSIBILITY. Possibility, then, includes anything that might or could happen within the context of the play's world. Hamlet at the beginning of Hamlet, for instance, could possibly return to school after his father's funeral. Why not? Forward motion in the plot is accomlished through the formulation of possibility into probability. As Hamlet unfolds and Hamlet has a conversation with his father's ghost, it becomes PROBABLE that he will remain in Denmark, at least until he has more information about the ghost and the circumstances of his father's death. Later in the play, as Hamlet verifies the ghost's story in various ways and as we see Hamlet's deep loyalty to his father and his desire for revenge emerging through the pattern of his actions, that line of probability is further refined to predict that Hamlet will remain in Denmark until his father's death is avenged.

In many plays, several major and often contradictory lines of probability exist through the play until the final moments. It appears probable that Claudius will eliminate Hamlet before he can avenge his father's death throughout much of the play, but as the action proceeds, it is also probable that Hamlet will succeed in killing Claudius. (By contrast, it is not probable that Ophelia will slay Hamlet in revenge for the death of Polonius, judging from her actions and what we infer about her character).

At the end of a play, one line of PROBABILITY is finally formulated into NECESSITY. Again in the case of Hamlet, the probability that the evenly matched wills and intentions of Hamlet and Claudius will produce a conflict that ends in the death of both characters becomes inevitable, and other lines of

probability are withdrawn. It is this final movement from probability to necessity that signals the end of the action (and of the play). From the vantage point of that final necessity, we can clearly see the causal chain of events that led to the outcome, producing the emotional and cognitive sense of closure that we feel at "the end" (we have probably all seen films of plays with "crummy endings" which failed to meet the causal, and therefore the emotional and cognitive criteria, of dramatic necessity).

The causal linkages of events, then, established the movement from possibility to probability to necessity and provides the play with its emotional shape. A story generation program like Tale-Spin is capable of creating one type of causal linkage which establishes probability and unifies the action of the story: the causality of goal-oriented behavior. The central character has a goal; there are obstacles to his goal which consist of the goals and traits of other characters (nested hierarchically); and the character ultimately either succeeds or fails. Here is a simple Tale-Spin example: