

IJCAI-95 Workshop Program  
Working Notes

## **Entertainment and AI/Alife**

Saturday, August 19, 1995

**Sponsored by the International Joint Conferences on Artificial  
Intelligence Inc. (IJCAI), American Association for Artificial  
Intelligence (AAAI), and the Canadian Society for Computational  
Studies of Intelligence (CSCSI)**

# Rethinking AI for Art and Entertainment

Phoebe Sengers

Department of Computer Science and  
Program in Literary and Cultural Theory  
Carnegie Mellon University  
Pittsburgh, PA 15213  
phoebe.sengers@cs.cmu.edu

AI applied to art and entertainment is in a whole new domain. It would only be natural for the science and engineering practices of AI to have to change — or at least expand — in response to the new challenges posed by this novel area. In this paper I will explore some of the fundamental changes in mindset that may have to be made in order to develop AI programs that do useful things in this new domain. Specifically, the mindset and assumptions in the ‘soft’ field of art and entertainment are very different from those of the ‘hard’ field of AI; my argument is that there is a culture clash between these two fields as they are now understood, and that making AI programs that are truly useful for people interested in creating entertaining social works will be made much easier by explicitly considering the differences between the fields and learning to develop a new kind of AI technology that is also at home in art and entertainment.

AI practitioners venturing into the world of art and entertainment are a little bit like the Europeans when they first came across America. A whole new field is open for exploration, where the technology we have has never been used and can point to all kinds of interesting new ways of living in this new territory. We could, if we choose to, proceed the same way the Europeans did: deal with the pesky natives by handing them technology that destroys their culture, clear them out of all the best areas, and start tearing up the environment to make it more amenable to the tools we already have. Aside from the obvious ethical implications, this strat-

egy has a major fault. In ignoring the cultures and technologies of the people who were already there, and in attempting to go boldly forward with tools that were not developed with this environment in mind, we would be in danger of destroying or at the very least failing to connect with precisely the things that attracted so many of us to this new area in the first place: not a quick buck, but the inherent interest of working in fields and with colleagues that engage people in positive, creative, fascinating, and somewhat alien ways every day. By imposing our own standards and worldview on this domain, rather than considering and perhaps adapting to those that are already in place, we are in danger of creating technology that is at least irrelevant to the domain and at worst destructive of it.

Obviously this story is a little overstated: I am not accusing AI researchers of plotting to take over art and put artists in intellectual backwater reservations. Yes, currently AI does have more funding and political backing than art, and more epistemological clout and technological gee-whiz-ness than either art or entertainment; still, a wholesale takeover of the standards of art seems impossible. But then, I don’t think the Westerners were plotting a dastardly takeover either; the destruction of native cultures and the environment came about gradually as a side-effect of a basic culture clash that the Europeans themselves did not fully understand. Cultural relativism was not a fashionable concept in the first few centuries of what we now know as the United States, and the consequence of

this is that Europeans took their own culture for granted. They were blind to the assumptions inherent in their culture and technology, and because of this had a hard time both having respect for the cultures and technologies that were already there and adapting their way of existence to their new environment. Instead, many chose to bulldoze that environment out of their way. The question is not who is doing what to whom, but how can we avoid making the same mistake the Europeans did, and end up in our enthusiasm about our technology missing the artistic point entirely?

Certainly many AI researchers in this area are already fascinated by the arts and the world of entertainment. Dialogue between the two sides is helpful in keeping the interaction between researchers and artists from devolving into the sorts of 'imperialism' with which Herb Simon characterizes his forays into literary criticism[2]. This dialogue has already begun, including the AAAI 1994 Believable Agents Symposium and 1995 Interactive Story Symposium, at which AI researchers, artists, and other interested people exchanged their views. I want to propose that there are a number of other things that we as AI researchers can do to make our attempts to interface with art and entertainment as easy on ourselves and as nondisturbing of 'native' traditions as possible. When we consider how to adapt AI we tend to spend most of our intellectual energy considering how the technology which we already have also applies to the new domain. An obvious, but not necessarily easy, adjustment to this technique to make it more domain-friendly is to also think about what kinds of (perhaps not-yet-developed) technology might naturally fit in to the new domain: what kinds of things would artists actually like to use? what kinds of things really fit in with their notions of what they are doing?

To do this it is not enough to take a straw poll of artists and creators of entertainment, because they probably understand the possibilities of AI technology as little as we understand their work. It really requires an understanding of the culture and possibilities of

both sides to be able to develop a technology from the old domain that fits comfortably and more-or-less naturally into the new domain. In particular, I am suggesting that we consider ways in which their cultures differ, and examine how AI technology is an integral part of the culture of AI. With this background we will be able to start considering alternative AI architectures that could better fit into the existing culture and goals of art and entertainment — or expand upon them in interesting ways. Here I will (nonexhaustively) list a number of ways in which the arts and entertainment traditionally differ from AI as an engineering discipline, along with some implications for these differences for AI architectures that want to work in these new fields.

### The role of the audience

Traditionally, designers of AI programs tended to build their programs by thinking about the requirements of intelligence rather abstractly, i.e. without much reference to the environment in which their programs will live. Behavior-based and other alternative forms of AI realized that thinking about that environment as well buys leverage in designing programs that can perform better by being less generally correct but more robust in their chosen domain. Even a superficial examination of arts and the humanities shows that their practitioners place great emphasis on another factor involved in the success of a particular program that has so far been to a large extent underrealized by AI. This factor is the audience of the program: those people for whom the program is designed and with whom we wish to communicate.

In science and engineering, we have the luxury of considering ourselves and our audience to a certain extent to be objective onlookers, so that we need concern ourselves only with what our programs are 'truly' capable of in the eyes of our colleagues. In the arts and entertainment, however, we are dealing with a more heterogeneous audience, and here the 'actual' capabilities of the program as defined and under-

stood by computer science are less interesting than the variety of ways in which the audience [as well as the designer] of the work may perceive it. A work will not be successful if it cannot communicate with the audience for which it is intended.

This means that for our audiences, deep structure is less important than surface appearances (although, of course, it is important to the extent that it shows up in the surface). This leads to difficult and, to us, unnatural changes in the standards by which programs may be judged. This is of course exemplified by *Eliza*, which, though laughable scientifically, still built up a cult following outside of computer scientists. In addition, thinking about the particular audience of the program as a demographic group may actually help to make the process of program design easier, just as limiting the design of an agent to the particular environment it will encounter simplifies matters. The creator can take advantage of a knowledge of what that group accepts, ignores, and finds interesting to decide how to build a system that will best speak to that particular audience. For AI researchers, whose audience has been other scientists, this has traditionally meant building goal-oriented, task-based tools; for other audiences, rationality may be unnecessary or even seem unnatural.

### The "point" of the work

In the engineering field of AI, an architecture or program is understood as an attempt to solve a particular problem. A well-designed program is useful for anyone interested in similar problems, and a program that included as part of its repertoire a long diatribe on politics or personal opinion would be considered rather strange. The opposite is generally true for the arts. A work of art — and many works of entertainment — are often considered to be a sort of communication, even if the artist him- or herself is long forgotten. A work of art tends to have a 'point' — including its own esthetics — which the audience should 'get,' rather than be a tool

to achieve something. Simon Penny, a CMU artist-cum-roboticist, has analyzed how the more utilitarian point of view that computer scientists hold has bled into the sorts of interactive art that they tend to create: he claims that the work itself is often intended to allow the user to reverse engineer the technology itself ("isn't that clever?"), rather than trying to communicate a point (social? political?) in which the creator believes outside of the technology[1]. This is also perhaps a side effect of the objective / subjective distinction: it could be that scientists are (probably correctly) afraid of making statements that might be seen as political in their art works for fear that the technology itself might appear to be tainted.

What difference does this make to AI architectures? Programs are generally designed in AI to be able to go out into the world on their own. The programming paradigm is that the user will take the architecture and design an agent or program that can do certain things. In an alternative paradigm that could be more amenable to the artistic point of view, architectures and programs would be seen as a method of communication between the creator and his or her audience. For example, in designing artificial characters it would become less important that an agent be able to be programmed to do particular behaviors and more important that the creator be able to specify the sequence or structure of behaviors at the level at which the audience will interpret them. We will need to stop talking so much about the behaviors of which the agent in and of itself is capable and more about the signs that the agents communicate to the audience, and the ways in which those signs can be manipulated by the builder.

### Messiness

Coming from science and engineering, clean architectures are of great value to us. Having a clean system with easily discernible parts allows other people to duplicate and build on our results; building a new system reduces to the problem of copying the old system and

changing some of the parameters. On the other hand, as anyone taking a stab at criticism knows, artworks are often terribly complicated and messy. In fact, part of their entertaining characteristic lies in the fact that you cannot easily exhaust them of meaning. A predictable and clean artwork is still an artwork, but it may be a rather dull one and it certainly encompasses only a small portion of the space of all possible artworks. In addition, the idea that someone would copy your work and then just twiddle a few parameters to create a new work is ludicrous. Each work is valuable because of and, to a large extent, unique in many of its details of composition. This leads to the possibility that clean architectures are not only unnecessary in art and entertainment, they may actually be counterproductive. It may be that a messy architecture that requires tremendous change, perhaps including hacking, from artwork to artwork is necessary in order to allow for the intricate, complicated detail that makes each artwork rich, interesting, and unique.

### The constraints imposed by the medium

In AI, we tend to think of computation as an infinitely elastic medium. If something can be done, then there is a way to get a computer to do it. Artificial Life, for instance, often makes the argument that life itself can be not just modeled by but incarnated into computational form. This faith is, to a certain extent, an asset — with it we have the perseverance to come up with ingenious solutions to otherwise seemingly unsolvable problems. This worldview clashes, however, with the view of arts, which understand each medium as affording different and unique possibilities for expression. A book and a film based on the book are not fruitfully thought of as the same thing; the specific properties of the media and the specific way in which the book is 'implemented' in the film make the film into a new and unique form of expression that should be judged more or less on its own terms.

This difference in worldview has var-

ious implications for AI-based arts and entertainment applications. The most obvious one (with which few AI researchers would dispute) is that electronic media can not be considered to supersede the old media; each category of media allows for particular kinds of creative expression, neither of which can indisputably be said to be 'better' than the other. Another uncontroversial effect is that when we are creating computational works we need to think of the possibilities inherent in the medium itself, rather than trying to translate directly from old media. Merely copying old media into computational format gives you a pale imitation of a medium, rather than exploiting the possibilities inherent in the machine.

A more subtle effect of this change in point of view comes from the notion of constraints. Artists and craftspeople who have explored a medium understand that media, while highly flexible, also impose constraint. Media do not allow for unencumbered 'implementation' of ideas, but rather give form and shape to ideas so that they become embodied in particular and highly specific ways. That embodiment may require subtle or drastic alterations of the initial idea: perhaps the original idea didn't really 'work,' or new ideas spring up from the interaction of creator and material. It would only be natural for this feeling of constraint — an understanding of the way the medium shapes the things one says and, in effect, prevents one from saying certain things — to extend to the computational medium. In this sense, working with artists or working as an artist, computer scientists can get an understanding of the ways in which the medium in which they work enforces the creation of particular styles of product, as well as some intuitive feel of the sorts of things that remain outside the possibility of computational implementation.

### Conclusion

AI and art and entertainment are at the moment wildly divergent fields. The conjunction of these two areas is not

just a case of tool meeting practice but involves the clash of two divergent worldviews. Rather than trying to impose the worldview of AI on that of art, I propose that the builders of AI think carefully about creating a new sort of AI technology with a hybrid worldview that will more easily fit in to the native goals and traditions of the field they are trying to affect. This will involve a rethinking of the ways in which AI technology flow out of and reinforce the worldview with which they are currently affiliated, and require both new kinds of architectures and new techniques for evaluating architectures. In order to build technology that is truly amenable to these new fields, scientists will need to do a good job of educating their colleagues that the current standards of AI are not set in stone. The only alternative is to try to impose the standards of AI on the arts and entertainment, thus failing miserably to join the field or eliminating the alien, yet beautiful, nature of that field forever. While the process of re-evaluating the standards and techniques of AI may be somewhat painful, it has a positive flip side: bringing AI technology to art allows for not only a new kind of art but a new kind of AI as well.

### Acknowledgments

This work was supported by the ONR under grant N00014-92-J-1298. The ideas in this paper were developed over a protracted period of discussion with the other members of the Oz project under Joseph Bates and with Simon Penny and Camilla Griggers. Nevertheless, all views are the sole responsibility of the author.

### References

- [1] Simon Penny. Personal communication.
- [2] Herbert A. Simon. Literary criticism: A cognitive approach. Manuscript copy.