Bill Clancey Interview by Hans Ulrich Obrist

Recorded at Museé d'Art Moderne de la Ville de Paris
November 2001

At the time, Clancey was a Research Scientist, Institute for Human and Machine Cognition, University of West Florida, Pensacola.

On leave at NASA/Ames Research Center, Moffett Field, CA
Chief Scientist, Human-Centered Computing

This is the original text, edited from a transcription.

Hans Ulrich Obrist: To begin with the beginning- I was introduced to your work by Israel Rosenfield and I started to read your books, and there seems to be a very complex position you have because of different activities. You have contributed a lot to the field of artificial intelligence and at the same time you are now one of the great experts for the Mars mission at NASA. I was wondering if you could tell me a little about how this all started in order to get a picture of your interdisciplinary activities.

Bill Clancey: I have always had a very broad interest in many fields. In high school I was reading philosophy, and I was following the space program in great detail. And I will never forget the day I discovered my father's psychology textbook from the 1930s and just started reading it and learning about it. And yet, computers even became accessible for programming, even as toys, in the 1960s. I knew by the time I went into college in 1970 that I would have something to do with computers, and maybe something to do with intelligence and psychology, and something to do with philosophy.

HUO: That was already there?

BC: That was there, and I knew my career would have something to do with writing, because I wrote compulsively, I enjoyed writing. The cognitive science field didn't exist yet. Even computer science didn't exist as a discipline when I went to college. In the early 70s someone showed me an article from Life magazine about a robot called Shakey from the Stanford Research Institute, and many of us looked at that and said 'This is it, we will go into artificial intelligence.' It was clear that there was a field that combined these interests of computers, psychology, and philosophy.

HUO: So, in itself, it is a very interdisciplinary field.

BC: Yes. People come at it from two different directions, as you may know. You have the physicists and logicians and mathematicians who believe that ideas and thought have a character that can be decomposed and then reassembled mechanically, and that the real essence of reasoning is logical deduction. Then you have the other side, which is the one that I was always coming from, that says, 'I am really interested in the mind.' I was also interested in psychiatry and psychoanalysis, so I was reading Jung and Freud in my college days as well, which was an inherent interest. As it developed by the 70s, there was no way to put all of this together. I took a straight computer science and AI degree. My dissertation work was to take an expert system, this model of medical diagnosis, and make a teaching program out of it. I always had the interest, which is one of the reasons why I went into AI, to build a program that we could interact with. That was the notion; some system that I could watch as a robot that could talk to me, that I could ask questions.

HUO: So that there is a feedback loop.

BC: Yes- the notion of HAL. What is interesting about HAL is that it is a being, that you can have some sense of an entity. Even as late as the 1970s we couldn't relate all of the disciplines of knowledge and learning and language, and by this I am referring to the work that was in linguistics or philosophy. I remember taking to one of my advisors, Terry Winograd, who was a very respected, well-known person...

HUO: Who is also in your book on artificial intelligence.

BC: Winograd was a classic MIT graduate. His work was in language. He was my advisor, and so I brought him these books. I brought Heidegger and Polanyi on tacit knowledge, and things that I had read. We couldn't connect them and he couldn't help me. Ironically, he was the one, when we fast-forward a little bit to the mid-80s, who came in contact with people doing hermeneutics, and began presenting a very different view on language interpretation and meaning, and rejecting the notion of messages as being packaged and transmitted; the information theory that all of our work had been based on.

HUO: Was there a link to Chomsky?

BC: Yes, but Chmosky's work, in a sense, as most people look at it and in the way that I usually think of it, is more in the

spirit of classical AI, because Chomsky's work on grammar was a strong confirmation and inspiration for people working on compilers. Most theories of natural language, and maybe AI researchers were taking a simplistic view of Chomsky, but essentially you could characterize all views of language and comprehension as being grammar based. Most cognitive models are descriptions of patterns that you find in behavior; a grammar is a fine description. We are not questioning that, but it doesn't have the generative capability of the human being. It is good for recognizing speech, and yes, you can use it to generate recognizable speech. But where is the meaning, where is the learning that would have to be behind it?

What I was going to say about Winograd is that in the mid-80s he comes out with a book called "Understanding Computers and Cognition." It was like an explosion in the middle of our field. It was an attack by someone who was extremely well respected from the halls of science at MIT, saying that the foundations of AI, and all of the researchers, couldn't be more wrong. He was going to give up the search for AI and begin work on the design of computer interfaces. He gave me the license, then, to go back to all of the other things that I knew about, and the main path that I followed to pursue these ideas was book reviews. I had started with some fairly traditional books in my field, and I found that you could take something that someone had written, and you were protected. You weren't proposing your own idea, you were promoting someone else's idea.

HUO: That is one of the things I found interesting about your sampling book on artificial intelligence, "Contemplating Minds." That is something in my field or fields, if it is art or architecture or literature, that I have not encountered before; that you use the book review as a medium. I was wondering if this is something very specific to the artificial intelligence field. A lot of these texts are like manifestos. The book review becomes like a modest manifesto. At the moment when we are no longer having immodest manifestos it is a modest manifesto. Would you agree with that?

BC: Absolutely. It is honest at the same time. It is a book and a person's work that, with full integrity, you believe in and that you want other people to read and you want other people to understand. It is a vehicle for promoting the ideas that you believe are right and that other people don't yet understand. Within a scholarly field, it is an acceptable way to introduce ideas. The book has already been published and the book review doesn't have to go through the same peer review process as well. There isn't this argumentation about 'are you right or wrong?'

It is your interpretation of the book and people find that interesting. We did a whole series of book reviews in the AI Journal.

It was Israel Rosenfield's book that I had discovered, "The Invention of Memory," just from browsing in a bookstore. That was a book that was pivotal for me because it provided more evidence for what I wanted to say from neuroscience and the history of science view, and a reinterpretation of those previous experiments and what was happening, for example, when a probe was put in the brain during surgery and someone remembered things. People had said it was evidence that there was localized memory happening. Israel provides an understanding of the different interpretations. This was exactly what my people needed to hear.

From there I did a couple of other book reviews and got involved with Gerald Eldelmen's work. Edelman was a Nobel Prize winner for his work on immunosuppression and he had applied his model, a kind of Darwinian model, to understanding how learning occurs in the brain as a competitive selection of neural maps. What I especially found useful from Edelman, first of all, was that model of learning, which was a non-descriptive view of learning based on interaction and experience; that the brain retains not a description or even a model of the outside world, but it becomes a system, an interactive dynamic system, that in motion, as it moves through a field, it perceives and is coordinating its perception with its action. That was very much in Edelman's theory. In a way it was good because his background was not very strong in psychology, so he didn't worry about the conceptual level and reasoning and all that we were doing. He focused on perception. That was a connection I needed. We needed to ground this new theory of knowledge to something that was connected to motion and perception. You may be familiar with people in Ecological Psychology like Gibson and Turvey, a group of people in the 80s. They started learning about AI and robotics and started to cross over. In my book "Situated Cognition," I tried to bring all of this together. I had a section on robots.

HUO: It is like a composite or a contraction.

BC: Yes. But also to show that there are all of these threads that have been going on for over one hundred years. My favorite was a paper by John Dewey of 1896, where he criticizes the Stimulus Response Theory of learning, which was just coming into formation. He says that it cannot be the case that the stimulus is somehow packaged and transferred by a linear process to an

action which is then packaged and sent out. He said it has to be a circuit that feeds back to a continuous perceiving to motion. I found this incredible because, then, for 80 years you have the Stimulus Response theories and behaviorism and, really, the grammatical theories of artificial intelligence; you have some input and then you have a rule and you make conclusions. You go from sensation to action. They all have this linear view. The very diagram that Dewey criticized in 1896 was being promoted in my field as our great discovery. He criticized it. He was so far ahead.

HUO: Dewey is also very important for my field. What is interesting is that Dewey has a relation to dogma and pragmatism and there was a whole idea about museums that developed. It was actually the museum of the future. It was a non-linear issue of how to cope with uncertainty and unpredictability and so on. It is an interesting link.

BC: I wasn't aware of that. I knew about it, certainly for education, so it makes sense to me. He was so strong. The other person that figures in my book who was key for me was Bartlett, the British psychologist. Bartlett wrote this book, "Remembering." I don't know how I came upon it. I am proud of one thing, which was making such a big deal with Cambridge University Press about Bartlett and "Remembering" that they reissued it right after "Situated Cognition" came out. I thought it was great.

HUO: And "Situated Cognition" was your book...

BC: Yes, it was before "Conceptual Coordination."

HUO: It was a book you did before "Contemplating Minds"?

BC: It was about the same time. I had conceived what would have been four volumes and "Situated Cognition" was the first volume. "Conceptual Coordination" was the second, and the works that are not yet out and are still in progress are more about consciousness. "Situated Cognition" was to tie together these streams of psychology, philosophy, robotics, and the ecological psychology. To me the most fascinating thing always, and as I said, this goes back to my days in high school, was that you could pick up a book in another field and you could say this is speaking to what I am interested in, even though this person is not a psychologist or whatever.

HUO: So it jumps?

BC: Yes. That was the wonderful thing about Israel Rosenfield and meeting him. I said, 'Here is someone who is an ally outside of my field, who is speaking for what I want to say and helping me and will support me." While inside your field the majority of people do not understand you, and not only that, but they say that once you were smart and now you have become crazy. We are learning a lot about the mind, I believe, by looking at the difficulty scientists have learning new ideas. I think one must not become cynical about it, and to understand why it is good that there is such conservatism in the disciplines, especially in universities in most departments. What I found, over and over again, and I could show you book after book, they take what I wrote, and also Rod Brooks at MIT in robotics, and greatly simplify it into a statement that is absurd. We can all agree that nobody would ever believe that. They will say that this is what Clancey and Brooks believe. Their favorite statement is that Clancey and Brooks believe there are no representations in the mind. They all repeat that.

HUO: Which is not true.

BC: It is absurd and no one could believe that.

HUO: What is the aspect in the work of you or Luc Steels, because he is also very close, and Brooks and Rosenfield, that general, mainstream fields in science do not want to understand? How would you define it for a non-scientist in order to understand it?

BC: The metaphor we use for understanding knowledge and memory is that knowledge is a substance, and memory is a place where knowledge is stored. We have a physical metaphor for how the brain works which is very easy to understand: the brain takes in information, it codes it, it stores it, it matches it, retrieves it, indexes it, composes it. It is like a manufacturing model. It is something in our everyday experience that we can understand. It happens to be how most computers work, and when you apply this metaphor for creating models of how people do geometry problem solving or medical diagnosis, you have fairly good models. Which is to say, if you take a human being like a physician, and you abstract the physician into a description of patients and a description of diagnoses, and you map the patients onto diseases, and then onto therapies, you can put that encoded as a grammatical rule structure into a program, and you say, 'Look, the physician is behaving exactly like the program.' But that is not medicine. The physician does not receive a description of the patient. A physician interacts with the patient and has to do a physical exam. And, by the

way, it is not the physician, but it was the nurse, who, at the reception station was interviewing the patient in advance. Having only built an "expert" program, you would know nothing about the practice of medicine. You have said nothing about what it means to be a physician. You have reduced medicine to these relationships between descriptions of patients and diseases. In doing this, we have lost track of what human beings are.

HUO: Israel once told me that he thinks of memory and knowledge in a dynamic view and what is so interesting is that, going back to Dewey and museums, it is the same thing; museums are places for storage of knowledge. Museums are not places for processes of unpredictable knowledge. Dewey and Dormer said, at the beginning of the 20th century, that museums should begin to cope with the unpredictable and with uncertainty and be complex, dynamic systems. I think there is a link. And this interview takes place in a museum.

BC: Where I see the connection is where Dewey said that the curriculum, your prescription of what a child must learn, is a map or a path that you might follow, which guides learning. It is not what you need to learn. It is the museum as a set of paths instead of the products that we want you to consume, or knowledge as a substance that you feed.... Like all those diagrams of knowledge pouring in the brain and so on. I think the other element that you are starting to hit on, where we have the connection with Luc Steels and his robotics work and a systems dynamic notion, is very simple: Every human act, speaking, and moving, is new. There is no such thing as simply taking a pattern and instantiating it as a robot works. We are always re-coordinating our sense of meaning. It is never just a rule like a pattern or like cookies being stamped out. That is what Bartlett said beautifully.

HUO: That was early, wasn't it?

BC: Yes, in 1932. He said memory consists of sequences or processes that have worked together in the past, literally and physically in some way, that are being dynamically re-coordinated. He talked about coordination, and Dewey talked about coordination. That is why I chose the title "Conceptual Coordination" for my book; to think of concepts as not physical, static things that are stored away in some network and retrieved, but they are part of a motion system, an emotion system. Of course, I am applying a metaphor in some way. I think it is a better metaphor--that we begin with a notion of physical coordination, and we try to understand perceptual motor

systems in animals and insects. Think of the conceptual systems as, simply, advanced physical coordination; a way of packaging sequences and dynamically recomposing them, and substituting them, and allowing a blending of sequences. That is what I explain in the book, looking at many examples where we can see in human behavior this blending and substitution. That is why I was fascinated by things like spelling mistakes and typing errors. People have done these studies and they have had models that they call "neural net" models that were very useful for me. What I tried to do was an analysis that is like a periodic table; to say, 'Okay, they list for me thirteen processes and I believe there are maybe two: composition and sequencing.' Now look at the variations of what could happen. It is partly my mathematical background, too, that I say the architecture must be something simple, and it is the right scientific start. We have years and years of descriptions of behavior, whether it is diagnosis or typing, reading, children's mathematics problem solving. We have excellent descriptions. I come in and say 'Let's view these as taxonomies.' The butterflies have now been laid out; can we, like Darwin, say what is the process that lead to this great variety in human behavior? The fundamental part, and again we come back to Luc and Israel and Bartlett, is to reject the substance model of memory. It is a rejection of the storage view of learning. The thing that is wonderful for me, which I convey in this book and the other, is that I didn't make this up. It is not my idea. We can find hundreds of people saying this.

HUO: It is a long story. There is also Whitehead talking about this.

BC: Yes. The traditional view of learning is that I have a learning experience where someone is teaching me, or I am reading a book, then I assimilate the learning, I store it and then, later on there is performance. In performance there is no learning occurring, it is simply my application of knowledge. Learning packages knowledge, it stores it, and performance retrieves and applies knowledge. Of course, and this comes back to the museum, this is our model of education; this is how we do school, this is how we do education. As we were talking about the other day, the notion of a traditional exhibit is: 'I am a specialist, I package up the presentation so that the museum goers have this encounter with knowledge that has been now been put into the ideal form, so it is accessible.' You digest it, you retain it and you move on to the next exhibit as opposed to learning by doing. "Situated Cognition" is part of Cambridge University Press's Learning by Doing series, which was Dewey's notion, again. I think the part that we are getting at today is

that there is a superficial view of learning by doing as well. You learn best on the job--anyone in any field. That is true, but it is at the neurological level that everything is an improvisation, everything is an adaptation.

HUO: That is why, for me, everything in neuroscience, like Rosenfield, has always been a great inspiration for exhibitions because it is this idea of complex dynamic systems with feedback loops. That leads me to another question before asking you about museums. I also wanted to ask you about museums because you mentioned this science museum the other day. Before coming to that, I wanted to ask you a little bit more about this dynamic notion and how far it could relate in terms of vision, fiction, and action. I just saw this exhibition which was in three parts and it started with vision, it went into fiction and then into action.

BC: When you say those words, the person who comes to my mind is Donald Schön. Donald Schön was very influential for a number of us. He developed Dewey's line of thinking. He was working with architects at MIT. He wrote a few books. He died somewhat young a few years ago. His books were about what he called the "reflective practitioner". He worked with musicians, architects, and people engaged in other forms of problem solving. What he was trying to show us was how we interact with materials as part of a creative act. He gives an example of someone designing a school. He says 'Let's look at this whole process,' and they put some marks down on a piece of paper, an expression. So, we say, now it is reified, it has become articulated in some sense. Now they perceive what they have put on the paper and they see relationships. Some are accidental or serendipitous; some were maybe implicit in their original thinking. He calls this "back talk". Now the representation out in the world gives you yet another idea, and you now change the drawing and it develops in this way. This notion that we put out into the world something that is based on a relationship that, maybe, is non-verbal, it is conceptual, it is emotional, and it is expressed physically. Then we interpret what we put into the world. This was central to what Rosenfield, Brooks, Steels, and I were trying to say about representations. The previous theory said that everything flowed from inside to outside, and that when it was time to perform, you produced words. Now we come to the fiction part and I think there is some connection. The essence that was so important for me to see was in language. The theories of speaking that are most dominant, and the best models that we have, claim that when I speak I have an intention, which is a pre-description in memory that is stored in some working memory

or what we call a "blackboard," and it is then mapped onto a grammatical expression by applying some rules about how the words must be ordered. Finally when my brain has assembled this sequence, it is like a train with cars and supplies. It then will be sent to my mouth and will come out. Always the claim is that before I utter any sentence, the entire sentence was constructed in advance on the internal blackboard. That is how most artificial intelligence models work. In contrast, the dynamic view of the environment, cognition, and representation says that, as I am speaking, I am hearing. I am doing simultaneous interpretation, if you think of it that way. I am adjusting the rest of my sentence.

HUO: It is a conversation piece.

BC: Schön calls it “conversation with materials.” It is a very nice phrase. So when you said vision, fiction, and action, in the Schön sense, this relationship between perception, speech and physical manipulation of the world is behavior. When we want to look at learning, when we want to look at creativity, when we want to look at any, so-called, expert performance, we need to see it as dynamic in that way. This was the part that for most artificial intelligence researchers was impossible to understand, because it claims that thinking is not something that happens before action, but is itself a behavior. As you see in my example with Schön, he said if you look at someone, who is creating something, it is a very tight feedback. It is happening inside and outside and it is happening with other people.

HUO: That is also a question of urbanism then if Schön leads to architecture. I have interviewed a lot of urbanists: Cedric Price and Yona Friedman, who, since the 60s, have defined what Price called the famous "non-plan" to get away from the urban pretension of a master plan because we cannot predict what any inhabitant of the city does next.

BC: Yes, in fact, Schön was Professor of Urban Studies and Education at MIT. One way that I have explained these ideas comes from an experience I had in a town called Seaside in Florida. This was one of the so-called "new towns"; an attempt to bring back to America the notion of a village, a place that was of a human scale. Part of the method is to look at successful patterns in communities that simply grew up naturally. You have very simple notions, such as you encourage parking on the street. Then when the traffic goes through, cars have to move slowly because it is too dangerous to go fast. Simple things like every house has a porch or a balcony to watch

over the street, because people like to look out and watch other people going by; there is no big expanse of grass in front of the house, you put the house right up on the road. These were patterns that architects had extracted from southern towns. At Seaside they produced, on one page, a simple type of matrix of different size homes, and they provided rules, like every house must have a picket fence but every single fence on a street will be different. You don't always know this walking down the street. It was part of the pattern that architects saw when they studied harmonious places.

HUO: How do you see self-organization out of this? The urbanists from the 60s have, obviously, talked a lot about self-organization. How do you see, in more general terms, issues like the conceptual coordination mind or the experience in time, self- organized systems?

BC: I look at it on two levels. We were just talking about the design of a town, so there is some self-organization among the inhabitants as they are producing something harmonious, something that fits. In the particular story of Seaside, I was trying to make clear how the rules were shaping the future. Simple rules are applying recursively on top of each other, and that leads to some of the new patterns that you see at Seaside.

The connection I wanted to make on the other end concerns self-organization in the mind, bringing me back to my book: I think the fundamental vision I have, it is just a direction, is the idea of a conceptual system. Every idea is related to the other ideas that we have. This was a point that Gregory Bateson made. He said that the conceptual system, our ideas, form a tautology. They are self-relating in some way. One of the obvious things we know about conceptual systems is that we all have multiple conceptual systems in our experience, and they are related to the different settings in which we behave. We don't experience inconsistency because we don't allow these systems to interact and they don't interfere. I gave in my book a very simple example, there was a famous person in the media whose name was Ron Goldman, and he was involved in a famous murder case. Now it so happens that I was sharing my office with a person who was also named Ron Goldman... Suddenly one day it occurred to me my office mate had the same name as this famous person in the media. Why did it take months to realize that? It seems so trivial, but it makes the point that we operate in these multiple conceptual worlds. To make progress in understanding cognition, we need to understand that.

HUO: And make bridges?

BC: Well, it is related to what I was saying about why it is difficult for scientists to accept a radically new idea. You can't assimilate something that undermines all of these other ideas that work so well for you and that are related to the tools that you have. We once tried to bring all of this together within my group; we were talking about the other people working on artificial intelligence, and we said, 'Well, what do you what them to do differently on Monday?' Here we were, we had all of this profound understanding, we had a new conceptual framework, and yet we had to be pragmatic, to realize that these other researchers were people who have jobs, and are teaching courses, and have these computer systems, and they are famous, and they say all of these things which you are now saying is nonsense. You can't expect that to simply change. You then become aware of these islands of scholarship and discourse. It relates, again, to the notion of a conceptual system, but now at a community level. When I was first a student in AI, and I saw the debates between MIT and Stanford and Carnegie Mellon. I thought it was a waste and inefficient. I had this idealism that- 'Why can't we come together and learn from each other and speak a common language?' When you are young you feel one year is too long for change to occur, and later you realize five or ten years is not too long. Later I realized it was good that the people at MIT were pursuing their idea of what they called "frames," and it is good that the people at Stanford were pursuing their idea of what they called "rules". I showed in a paper in 1985, called “Heuristic Classification,” that at another level of analysis, these formal languages were the same. The conceptual structures that they represented were identical. There were different emphases--a frame system had something that resembled rules, and a rule system had something that resembled frames.

HUO: What was the publication?

BC: “Heuristic Classification.” It has been in some collections of AI research. It shows the style of my work and what I was trying to do. Everyone was focusing on these languages for encoding knowledge. I said that was fine, and now that we have collected these butterflies, what have we learned about knowledge? Well, look at this pattern. And again, the different researchers and their views gave me a way to put these together.

Let's look, for example, at our medical expert systems. In these computer programs, we have a classification of patients, by which we describe a patient as someone who has lived in a

certain way with a body with certain characteristics and now experiencing these symptoms. And then we have another classification of diseases, and then we have a classification of therapies. What these "knowledge-based" programs are doing is creating a model: I have a general classification of patients, where does this patient in front of me fit? I have this and this combination; someone of this age and experience, with this type of job, showing these types of symptoms, showing a fever and so on. Then I map that and say 'Where is the association to diseases?' And here is where the notion of heuristic is--it is not certainty, it is not a logical proof, it is associational, based on statistics or weak causal arguments. Then once I have reclassified this person as having a disease, I now say, 'Okay how should we treat that patient?' We need to make an association to the classification of therapies. Again, there is no certainty depending on the disease. We could say this kind of regime applied in this way, given the particulars of your case, this is what I recommend.

I called this uncertain mapping from one classification to another "heuristic classification," It was a simple mathematical hypothesis to say this is probably a good high-level view of the nature of professional expertise. We classify from the immediate information we have of the people and the situation, and we map to an abstraction, another classification and then we map to action. 'If this is my model of what is going on, then what should I do?'

I was influenced by Jerome Brunner, who had done basic psychological studies of classification. I was reading his work on classification, where they would look at single classifications and usually simple ones, maybe they were patterns of blocks or cards of some sort. My field was all about heuristics. I just put the two fields together and I said, 'I think what my field has done is to show that, given multiple classifications, experts know how to relate them.'

For Schön and Bruner, the idea of narrative is a central part of intelligent behavior and expertise. I mentioned before Schön's phrases "conversation with materials" and "back talk," He made me realize that if we want to understand the brain we have to see that the capability we have is fundamentally to construct narratives. That is very appealing from an evolution of cognition perspective because it is easy for us to understand that story telling was inherent in human society before writing. But the fact that I am getting at is that at a neurological level we need support for this process of "sequentializing." Why do we sequentialize? Why does sequentializing give us a feeling

that we can grasp the world? I think that is the essence. What Schön helped me to understand is that before there can be a causal story, there has to be a temporal story.

HUO: So it is all about time.

BC: Yes, exactly. One should then have a notion of a sequence of events and you can tell a story to yourself of what you have experienced or what you have seen or heard. Now you can attribute causality to the various events and you can say the causality isn't necessarily linear.

HUO: It is so interesting in relation to time because it leads a bit to the discussion at breakfast this morning about Mars, because I wanted to ask you a little about Mars. It may be abrupt now, but there seems to be a link. When you talk about this time notion and once, we said, in ten years we will go there. Is this somehow related?

BC: I think in a general sense. Now we are talking about large organizations.

HUO: It goes from micro to macro.

BC: Now we are at the notion of dedication and motivation in society today. How do those of us who would like to galvanize society, and lead, and get people to leave behind some of their mundane worries, give them some grand aspiration? We have to set goals that can be in their lifetimes so they can clearly see that their children will participate in this. You need that. Instead, if I told you, let us start building a cathedral, which would take one thousand years, you would say that is not possible. I suppose human nature must have changed somewhere along the line. To me this is a great mystery. Being in Paris and Europe, we have to ask ourselves what was the society that could be reaching for the stars and the sky, and building these cathedrals, and it was okay for it to take hundreds of years. They knew from experience it would take a very long time, and yet they proceeded and were able to organize themselves in a way that many scientists have trouble doing today. I don't know if it is the government funding or maybe it is the dynamics of the world, because we are not isolated..

HUO: Is it a utopia? One of the questions I was asking myself is the question of the utopia. I am making, at the moment, a big research on utopia and utopias always have to do with bridges. If you look at Thomas Moore in the second edition of his Utopia book from the 15th century, has a bridge to the land

from the island where the utopia is. Very often if one talks about artificial intelligence, one things of it, maybe, as being a utopia and at the same time, very often, if one thinks about Mars, one thinks, also, of utopia. I was wondering if there is a link there?

BC: You are totally right. Having attended three Mars Society conferences and listening to people... even though people are interested in life on Mars I think the geologists have an interest, but most of us can't get too excited about the geology of Mars. The only interest for most people in the geology of Mars is when they start talking about floods. They say there was all of this water, and there are pictures I can show you, a rendition of what an ocean would have looked like on Mars. That image suggests—we get now to a utopia—that Mars could be like Earth. The idea is that, and this is very speculative because maybe in two years we will figure out that this is nonsense,...

<<TAPE CUTS>>

...that is the exciting part of our times. We are sending probes, right at this moment, to Mars to find out these answers. Some of these things we are going to know in February. It is astounding. But utopia... I think the vision that attracts people throughout the Mars society...

<<NUMEROUS BREAKS IN TAPE 50:00min. to 1hr. 5:00min.>>

Historically people would go on an expedition, and many of us have been reading about the search for the Northwest Passage in northern Canada, because that is the area where we are doing our work. As you may know, they would get on a ship and they would be gone for months and many of them were never heard from again, or they would come back a year or two later because they got stuck and there is no communication. Today through electronics, we will be in communication with the crew on Mars. This is something that fascinates me. How should we establish the communication between the crew, the astronauts on Mars, and the rest of the Earth? The traditional view with NASA is one line of communication, and so everything must go through Mission Control, or perhaps better, "Mission Support." There is some truth to that, but the internet and electronic media will allow us to share—not in real time but five or twenty minutes delay—with the rest of earth what is happening on Mars. It seems inevitable that Mission Support will have to provide this relationship to the public and to other scientists, artists, and whomever.

HUO: Which is what we do today at the museum with Melik's project, where it is a very related issue where you will be on Mars in Melik's space and you will communicate with the public.

BC: At a time delay.

HUO: This participatory artwork of Melik hits this very precisely.[1]

BC: Absolutely. The challenge for me, and now you are very precisely down to why NASA pays me and what is the research. It is that we have to help conceptualize this communication between Earth and Mars. We are subject to very practical constraints, like you can't have six people communicating with six billion people directly. There must be some kind of filtering and privileged access. On the other hand, you can't have such an expensive expedition that is so important occur in isolation. Since the expedition on Mars would last about 500 days on the surface, there is plenty of time to seek help. We can have an international symposium where we bring together people, and we say "Here is what is happening on Mars today and we would like your reaction so that we can decide what we should do next month on the expedition." It is very different from the moon, where you are there for two or three days at the most, and there is no time for this participation. Now you can have the astronauts prepare some video, send it by email, put it on the Internet, and show it in the museum.

HUO: That is an interesting relation to filmmaking because you are also shooting a lot of videos of the expedition. I followed your expedition this summer by email and the website. You mentioned before that artificial intelligence is a very interdisciplinary field, but I think Mars Studies, if one could call it that, which I became aware of was an interdisciplinary field also. You mentioned Cameron, giving lectures and being interested, and you have artists and science fiction writers. I wanted to ask you a little bit about the link to literature, to cinema, to science fiction, your dialogs with Cameron and this whole field as an interdisciplinary field, and if it is true that it is as interdisciplinary as artificial intelligence. It is many questions in one.

BC: I think, for all of us, I would begin with the media. Video and digital photography have made accessible to all of us,

---

[1] Refers to Melik Ohanian's "Mars Room" exhibit as part of the Traversées Exposition at the Musée d'Art Moderne de la Ville de Paris, France, November 2001—a room, covered with images of Mars' surface from the Viking Mission 1976, in which Ohanian conducted interviews.

to scientists and the public, records that are so compelling. This is what Cameron is relating to. Someone with the proper filmmaking background, with a sense of the story and the image, could convey to the public something that could be entertaining and yet educational and motivating. It is a personal interest of mine. We have had a series of movies about Mars, and they have been more or less successful, but they are based on fantasy, they are not based on what we are going to see when we go to Mars and they are not about why we go to Mars. Can we do something else for the public, where the actual image is even more compelling than the fantasy of something that doesn't exist?

HUO: How is this concretely happening? Can you give me a few concrete examples?

BC: Cameron is producing an IMAX video and a television series. The point you were getting to is that there is a real collaboration occurring. You have Zubrin, a rocket scientist, who has written a science fiction novel to be proposed as a script for the movie. You have Cameron, who is the cinematographer and filmmaker, designing for NASA habitats that people could live in on Mars, as part of his development of movie props. He develops a rover, a kind of house van that he is going to give to the Mars society, so we can use it in the Arctic as part of our expedition research. You have a complete interleaving of the writing of fiction, the production of movies, the development of props of designs, and of the research. There are no barriers. The people just have a common interest and they also have some common talents. You have the filmmaker as an engineer, and you have the engineer who actually writes a pretty good novel.

On a different topic, I wanted to make sure I didn't give a mis-impression about artificial intelligence, because I think for the missions on Mars, especially the early ones, to be successful, we need better-automated systems than we have today. There are a lot of people who have focused on robotics, the machine that will do some exploration independently. But the automated machinery that is most important is the life support system. It is the system that can watch the electrical power system, the water recycling system, the air recycling system, and can monitor this very complex process, adjust it in real time and communicate with people when a diagnosis is necessary. What we are having to do, to support a Mars crew of only six people, is to move into software the capability to monitor, control and diagnose these complex life support systems. That is, to me, the obvious home and role for artificial

intelligence. I prefer the phrase "model-based software." It is software that has a model of some kind inside, so that it can understand, so to speak, and adjust controls appropriately. I just wanted to make clear that I think the techniques that we have developed have an essential role to play.

HUO: So there is not a rupture with your work from before.

BC: Not at all. I'm in the Computational Sciences Division at NASA, and we are called the Human-Centered Computing group; the other group is called Automated Reasoning and we are part of one picture. They are trying to develop this automated software and we are trying to say, 'Well, how could this software understand what the people need?" What if you put inside the software a simulation of what the people are doing inside the Mars habitat? The software has a model of what the crew will need tomorrow. We have learned that there are some very simple constraints. For example, you have to burn some of the waste, and that requires oxygen. There is a process that is storing oxygen to be used for the incinerator. There are other things the scientists are doing that require oxygen. They are doing some experiments of some sort, obviously they are breathing. Maybe they are going to take a supply of oxygen and put it inside the rover, so the day after tomorrow they can make this long excursion over the horizon and will be gone for a couple of days. In this case, the AI system, if it knows the plans people have and how they are going to use these resources, it could decide, well, rather than using the oxygen today to burn the waste, I should wait until Thursday on Mars, and then we can do it then. It is a very simple idea, but when you are in a limited, closed system, if these things aren't taken into account, it is chaos.

HUO: There was one last thing. We spoke about your work in artificial intelligence and we spoke about your book, we spoke about you work now for NASA and Mars and I was wondering what, after having realized all of these things, what are your un-built roads, what are your unrealized projects?

BC: They are there, yes. One thing you and I spoke about, briefly, the other day is to do a book on Mars studies that would convey my experience and make it accessible to the public. That is number one on my list. It would convey current ideas about situated cognition, but it would especially show this idea of observing the expedition, of being a member of the group you are trying to understand. I want to convey how you learn from people in the Arctic something that is useful on Mars. The Arctic is not Mars, so there must be some kind of a mapping. So

you can show that even though some feature is different, the rest is the same, and so I can make a generalization. That fascinates me a lot. The gravity on Mars is different from the Arctic, so does that make everything that you learn in the Arctic invalid?

The other project is writing about consciousness. I have some unpublished manuscripts, and I am trying to decide if I want to publish in specialized journals and get some good feedback from people who work in some of these areas. I have writings about visualization, and autism, and the evolution of consciousness, and most of it is about dreaming. That is something I have been working on for 35 years since I was in high school.

HUO: That is why we were so happy to have your text for the Dreams book.

BC: Yes, you asked for a dream and I wrote about Mars. It was exciting to be part of your project.