

Research-Technology Management

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/urtm20

Scientific Design of Tools for Amplifying Intelligence

An Interview with Bill Clancey*Jim Euchner talks with Bill Clancey about the design of technology to reflect the social system of work.*

Bill Clancey & Jim Euchner

To cite this article: Bill Clancey & Jim Euchner (2021) Scientific Design of Tools for Amplifying Intelligence, Research-Technology Management, 64:2, 11-17, DOI: 10.1080/08956308.2021.1868912

To link to this article: https://doi.org/10.1080/08956308.2021.1868912

Published online: 01 Mar 2021.

Submit your article to this journal

View related articles

View Crossmark data

Full Terms & Conditions of access and use can be found at
https://www.tandfonline.com/action/journalInformation?journalCode=urtm20

Scientific Design of Tools for Amplifying Intelligence

An Interview with Bill Clancey

Jim Euchner talks with Bill Clancey about the design of technology to reflect the social system of work.

Bill Clancey and Jim Euchner

Jim Euchner [JE]: Bill, you started your career by building one of the very earliest tutoring systems using artificial intelligence (AI) methods. You recognized back then that the computer needed to have some understanding of the person it was helping. Your later work generalizes that to design of computer systems to amplify intelligence. Can you talk about how that early work evolved into your current focus?

Bill Clancey [BC]: My doctorate project was the first to use an expert system (MYCIN, a program that diagnosed infectious diseases) for instruction. The most important thing to know about expert systems—what we called symbolic AI—is that these programs worked by interpreting models of thought processes. I think of the early AI researchers as being like alchemists whose wild dreams led to chemistry. In what we called "knowledge bases," we were attempting to replicate human intelligence, but we wound up making a major contribution to science and engineering. We developed new computational methods that were able to model processes and systems non-numerically in networks. These included ways of representing taxonomies, causality, and events. In an instructional program, the models include the subject matter (disease processes), reasoning (the diagnostic process), and instruction (a pedagogical process). All of my subsequent work has involved explaining this idea of non-numeric modeling and applying it to designing and building automated systems.

JE: Why do you think it is necessary to create a new design paradigm for AI systems?

BC: Many automated systems have failed us. I think the real challenge is with megaprojects—projects that cost hundreds of millions or billions of dollars. When you really dig into them, you find that when these projects fail, the root cause is related to design: they are R&D failures, but they manifest as operational failures. Someone is driving a Tesla on autopilot, and they crash into a 50-foot-wide refrigerator truck. The poor guy who is at the controls when the thing had a catastrophic failure is invariably blamed. For sure, the driver's trust in the system was misguided, but why didn't the engineers know their system could view a huge white truck as being part of the sky? The failure of the Boeing 737 Max 8 MCAS system is similar. Originally, the crashes were blamed on foreign pilots who did not have the right level of training, but the fault originated with the people who designed and managed the system. These are R&D errors; the failures occurred at the time of design.

What do I mean by that? The system's design simply did not account for the interaction between people and the automated system, what we call the work system. Sometimes, such an error doesn't cost lives, but it does cost a lot of money. Let me give an example. The planning for the Space Shuttle in the 1970s and into the 1980s assumed, based on detailed simulations, that the turnaround time between launches would be 10 days to two weeks. This assumption led to a business model that included commercial launches that would justify the shuttle fleet's costs. In practice, the turnaround time turned out to be at best one to three months, and sometimes worse than that. Why?

Bill Clancey is a senior research scientist at the Florida Institute for Human and Machine Cognition. He earned a PhD in computer science at Stanford University and a BA in mathematical sciences at Rice University. He has pioneered artificial intelligence applications for medicine, education, finance, robotics, and spaceflight systems. He codeveloped ethnographic methods for modeling work systems at the Institute for Research on Learning and agent systems for the International Space Station and Mars exploration at NASA Ames Research Center. He is a Fellow of the American College of Medical Informatics, Association for Psychological Science, Association for Advancement of AI, and the National Academy of Inventors. His Information Classification: General book, *Working on Mars: Voyages of Scientific Discovery with the Mars Exploration Rovers*, received the AIAA 2014 Gardner-Lasser Aerospace History Literature Award. wclancey@ihmc.org

Jim Euchner is editor-in-chief of Research-Technology Management and Honorary Professor at Aston University (UK). He previously held senior management positions in innovation leadership at Goodyear Tire and Rubber Company, Pitney Bowes, and Bell Atlantic. He holds BS and MS degrees in mechanical and aerospace engineering from Cornell and Princeton Universities, respectively, and an MBA from Southern Methodist University. euchner@iriweb.org

DOI: 10.1080/08956308.2021.1868912
Copyright © 2020, Innovation Research Interchange.
Published by Taylor & Francis. All rights reserved.



Bill Clancey is a senior research scientist at the Florida Institute for Human and Machine Cognition.

We find that engineers often do not take into account how the people do their work when they simulate the work processes. Instead, they abstract and idealize work in terms of tasks—they model *what* is to be accomplished and not *how* the work will be done—the intricacies of what people actually do. When you have a purely functional view of work, you tend to take as given what the inputs are and leave out all the reasoning, negotiation, etc., that people do. We saw this in early versions of tools used by planetary scientists and engineers to program the rovers on Mars.

JE: The model falls apart because engineers haven't taken into account the way the world actually evolves, the logistics and workarounds.

BC: Yes. Task simulations have a notion of people, for example, but the people are often modeled just like machines; they process stuff as it comes along, like on an assembly line. In this worldview, the people just have to be trained to do their task, and everything will function smoothly. To a large extent, machines and people are considered interchangeable. When this is the mindset, people ask questions like those I heard at NASA: "When does the human become the tool of choice for solar system exploration?" People are tools! It's a very short step from describing what people do in terms of functions to calling them tools.

There are many examples of major R&D failures that occur because the underlying notions of work, technology, and people are flawed. The problems we are trying to solve are sociotechnical problems, and the tools and mindsets we are using are technical tools—without the personal and the social parts taken into account. Consider what happened when Uber placed a self-driving car in Tempe, Arizona, where the people had no place to walk. The Uber vehicle hit a pedestrian crossing at night in a place where there ought to have been a crosswalk. Existing problems in the world—problems that people were able to adapt to in the past—are exacerbated when we take an already complicated or even dysfunctional system and add advanced automation to the mix.

JE: Why is thinking that incorporates the way people actually live and work so hard for us? What about work and technology do we get wrong?

BC: Designers tend to describe work in terms of tasks because it's clean and simple to do so. These models describe how we'd like the world to operate. There's the work to be done, which is a kind of specification, and there may be a set of procedures that people are supposed to follow. Then there is the technology that performs or assists with a task. People are doing tasks and technology is automating other tasks; that's what the designs look like.

That's not the way that I view work or design. I adopt the ethnographer's or social scientist's view: work is a set of activities, a set of behaviors; and it's what people actually do—where, with whom, and with what tools—hour by hour, minute by minute.

The idea that you could design separately the work people do and the technology that supports the work is completely nonsensical to me. If you introduce new technology, you're changing the behavior of the people using the technology, so you've also changed the work. It isn't a matter of whether it's a good idea to design new practices at the same time as you are designing the technology; you are always doing both, and it's a good idea to think about that deliberately.

There is a mind bug of the engineer that says, "I'm going to make things better by introducing technology to improve efficiency; let me identify the functions that I can automate." They come up with an automated system, and then they say, "I need to put an interface on this system, and then I'll train the people to use the interface." I saw this at NASA, and we both saw the same thing at NYNEX. It is based on a linear notion of work system design: tasks–technology–interface–training. And it is based on two big fallacies.

The first fallacy is that people are infinitely malleable and that you don't have to worry about what they're already doing—their adapted practices—or the time and place that they are doing the work. You may not even know whether the functions you are automating are desirable to automate, or what will happen when you try to do so.

But how the world changes is not linear! Work practices are organic. A work system cannot be completely predicted or created by a design we specify; it has to develop incrementally. There are always iterative refinements. In reality, the automation of work is an empirical experiment that involves technology design as part of work design, and includes revised roles, methods, schedules, and so on.

JE: You focus on an activity view of work, which captures its emergent nature. What is an activity, and how is an activity-oriented view different from a task-oriented view of work?

BC: Well, first I want to be clear: a task description is always useful; it just isn't complete. We always need to formulate requirements—what a person or group is supposed to accomplish, their objectives. But a task/functional view of work doesn't specify *how* anything is done. People may follow procedures, but that doesn't capture how they accomplish each step. The procedures are by necessity incomplete.

This became clear to me when we were modeling business processes at NYNEX: a task model would specify that a fax was sent from one office to another, conveying a job order. An activity model would specify the location of the fax machines, who used the machines, how long it took to get to the machines and use them, when the faxing occurred during the day, where the recipients carried the faxes, whether anyone was informed as a result of the fax, and so on.

Another realization is that people can engage in activities based on motives and effects that are not functions—such as listening to music on the radio while driving. A lot of life is like that. If you're watching *The Crown* on Netflix, what is your task at that moment? You have no task (and might not be happy to be interrupted with one), but you do have an activity.

The most important point is that we are always engaged in activities, even when we are sleeping. I realized this most vividly when I was in the Arctic doing an ethnographic study during a Mars mission simulation. We were in our space suits and preparing to go out of the habitat. We had to stand in an airlock for some time and pretend that it was being decompressed. We were just standing there, just waiting, and I realized that this was an activity! A task-oriented view would say that you are accomplishing nothing; your suit and airlock is doing all of the work. But you need to be standing in a certain place for a certain time, and that's your activity. I've described other examples like this in my papers.

There's another key idea: it is within activities that we discover how things are going to be done, and whose knowledge will be called into play. Tasks are formulated in meetings, reading email, or maybe speaking with a colleague at lunch. Those are all activities. The activities of people are a kind of invisible fabric through which tasks are defined and doled out, and you won't find them in job descriptions. The problem-solving that they do is not like problem-solving in classic psychology experiments. It takes place in the real world. As we saw in NASA's Challenger and Columbia disasters, *who is allowed to participate* affects whether problems are even acknowledged, and when they are, what methods are allowed to be used to address them. If roles and responsibilities are confused, then the best technologies will be of no help.

JE: What comes first? How do you start the design process?

BC: You start by trying to understand the existing practices. Ideally, you do this in collaboration with the people doing the work—in other words, through participatory design. When NASA Ames researchers were assisting the Jet Propolusion Lab (JPL) in designing science operations for the Mars Exploration Rover, we used this approach. When thinking about technology design, the social scientists at the Institute for Research on Learning and NYNEX taught me to ask, "How are the people already succeeding?" Whatever work you are trying to improve, people are already adapting and learning to get it done! The questions to ask are: About what? With whom? Where?

I would ask people doing the work, "What accounts for the quality of the work that you are doing?" That's an open-ended question, and it has a helpful, positive spin. It's not based on the notion coming out of psychology that you have to mitigate "the human factor" and somehow eliminate the foibles and vulnerabilities of people. This is a different way of analyzing that assumes that these people are smart; they're succeeding; they have expertise. They have ways of coordinating people, technology, and their environment and have formed a work system that has evolved over time and is succeeding. So let's understand that. When I do this work, I start by formulating a dependency tree to identity what part of the existing practice contributes to the success.

For example, one of our conclusions about the Mars Exploration Rover scientists was that they are both systematic and opportunistic. The question for a tool developer to ask is, "What tools and practices—what routines, procedures, policies, and roles, etc.—make those two parts of exploration possible and contribute to the quality of the work?"

JE: You've developed a language for simulating the activities of people working with machines and with each other. Can you tell us about it?

BC: It's called Brahms, and it is a language for describing and simulating activities. Agents engage in activities, and agents might be people, robotic systems, or software programs. Activities are represented by a set of rules, which we call "workframes." Objects such as vehicles and instruments are modeled in a similar way. The behaviors of all agents are located in time and space, so we model where agents and objects are located and how they move in a simulated environment, which might be places or paths in buildings, on highways, in the airspace, or on the lunar surface.

We model at the level of detail that is important for how the model will be used. One agent, for example, can represent a whole group of people if that serves the purposes of the model. When we modeled the Mission Control Center back room at NASA for the International Space Station (ISS), for example, we simulated in detail how particular individuals managed file transfers between Earth and the ISS computers because we were aiming to automate that, but we didn't simulate the shift handoff in any detail. One of the fundamental notions of Brahms modeling is that the model needs to describe *behaviors*—where people are, what

technology they're using, and how long it takes. We model things like looking for an object in the environment, for example—finding a tool, locating reading manuals, reading text on a screen, or calling someone for information.

JE: The hard part is realistically modelling the behaviors of agents that are not computer systems. You cannot capture the messiness of actual work in a task model, because it misses all the spontaneous interaction and all the emergent behavior.

BC: As the researchers at the Institute for Research on Learning emphasized in the early 1990s, work is situated, dynamic, and socially constructed. It happens in a specific place, with all that implies; it evolves as the world changes; and the work is inherently social. People have roles and follow norms, and these norms help to define our understanding of who we are, what we are supposed to be doing, and how other people will assess what we are doing, even when we are working alone.

As for modeling spontaneous interaction, the Brahms framework is object-oriented. This allows us to simulate emergent processes. Each of the individual agents and objects—whether a person, a camera, an AI robot, or a vehicle—is an independent process in the model, just as they are in the world. A Brahms simulation maintains a model of each object or agent and its local behaviors, which depend on its activities, its environment, and the situation that other agents present to it. We model agents so that they can interact in space and time with other objects and agents. They are all running simultaneously, which is where the possibility of unanticipated behaviors emerges. Unlike workflow models, the sequence of activities in Brahms is not foreordained, even if the goals are.

Here's an example of what happens when you don't consider the interaction between a person and a complex machine in the real world. It's an analysis of Boeing's 737 Max 8 by Captain Sullenberger, who is well-known for landing an out-of-control plane on the Hudson River. If you haven't read his testimony to Congress in June 2019, I really recommend it. There is a passage in the testimony where he talks about using a simulator of the 737 Max 8 flight system. Even though he knew what the Maneuvering Characteristics Augmentation System was going to do, he didn't have enough time to recover and get the horizontal flaps back under control. He didn't have time to flip the switches he knew were necessary to disable the automated system—which was doing exactly the wrong thing, by the way, working against his actions, even as he was flying the plane in manual mode. He knew what was about to happen; he knew what he was supposed to do to address it; and he still couldn't do it! There wasn't enough time. This wasn't simply a matter of training; it was a matter of a failure to design the work and technology together in a way that could be successful in practice.

How did that happen? Did anybody question whether the timing assumptions used to verify the model were accurate? All we know for sure is that they didn't properly validate their model of flight operations. The design of a system like this has to be empirical. You've got to simulate people's practices to make predictions about what they can and might do, and then you need to try the new technology with people. Task-oriented simulations don't capture sufficiently the real-world time-space interactions that may occur.

An agent-based model that is built around activities is especially powerful for simulating work that is asynchronous and distributed—work that involves different aspects that are in process at different times and in different places. A proper simulation of aircraft automation has to model separately what the people are doing and what the control systems are doing. The interactions can't be simply assumed; the timing and missed connections need to emerge as behaviors of the simulation itself—as they're output, as what we observe happening.

As another example, consider the case of astronauts on Mars. Mars is rotating, with its own daylight and nighttime periods. We will have satellites in orbit around Mars that are only able to communicate with the astronauts at certain times. We will have people back on Earth who are trying to supervise and monitor what's happening on Mars, perhaps distributed in their own time zones. We have all of these moving parts! Task models presuppose the flows: the flows of information, the flows of control, the flows of data. But in reality, all of these flows are emergent, they are sequenced and located according to circumstances as the behaviors of agents affect one another. An object-oriented simulation of the total system of activities can capture that.

JE: That is true in a dramatic way for Mars, but it's also true for someone fixing a telephone line here on earth.

BC: Exactly! It is true whenever you have work that is distributed and asynchronous.

JE: To just fix a telephone line, there is often someone in the field, someone in the central office, someone in a back office, a database that might be right or not right, a tool that might work or not work, a customer who is putting pressure on you, metrics you have to meet. All of those things are mediating what you attend to and the decision you make at any point in time.

BC: Absolutely, and you know it was the NYNEX example that you gave to us back in 1992 that made this clear and that led us to develop Brahms as a distributed framework of objects and agents just "doing their thing." In Brahms, every agent has a set of beliefs about the state of the world—an internal model based on inferences, communications, text read on a screen, or something that one or another agent said. Beliefs change when agents perceive the world while they are acting in it. Working with Maarten Sierhuis, a computer scientist at NYNEX S&T, and the social scientists, we knew that we needed to simulate activities such as reading,

speaking, looking, moving, and so on, if we were going to successfully simulate how the work is done.

I started by thinking about each agent as being like an expert system. It has its own set of rules that guide its behavior. But the model is more than inferences, judgments, and changing beliefs; it includes how the activities of agents structure perceiving, communicating, moving, and the manipulation of displays, tools, and materials. The actions of multiple agents and objects transacting in the world influence each other in unpredictable ways. That brought us to the object-oriented framework and to the idea of a simulated world.

JE: You allow what happens to emerge. If somebody is paying attention to something else when an event happens, then you'll simulate the same failure that you might get in the real world if someone gets distracted. You're helping to make the invisible visible.

I think that the invisibility of activities is one reason that technology to support work is so often badly designed. It is easy to just focus on automation and efficiency when the effects on people and on work quality are invisible. I worry about this in the age of AI.

BC: I have been concerned, too, which is why I have been studying megaproject R&D failures. But I am encouraged today. There is a movement towards a broader view of AI—not as a standalone agent but as more of an enabler, a means of amplifying human intelligence. That's a productive shift.

The emerging consensus is that the work systems we are designing involve groups of people cooperating and collaborating, facilitated by AI systems embedded in their activities. How does all of that work together? In a special issue of the *AI Magazine* (Volume 36, 2015), Ken Ford's editorial referred to "cognitive orthoses"—basically "human-centered AI." It's the same idea that Ed Feigenbaum promoted at Stanford in the 1970s—expert systems would be amplifiers of expertise; they would build on the expertise of people and allow them to solve more difficult problems more quickly and more accurately.

JE: I think that making that vision a reality is the crux of the work you're doing, Bill! People have been saying for decades that computers should support people and not the other way around, but very little toward that end has happened in practice. How do you design systems for collaboration if the things that will make collaborative work effective are invisible? You have to be able to model and understand the interactive and emergent nature of the system to do good design. It's not just about human factors but about work design. That's what you need a new kind of modelling for.

BC: You're right. The focus has to be on the people, and that requires a different approach to design. There are two big things that I want to put on the table. First, I'm promoting a scientific design approach. It's the idea that the design of a system is empirical; it's based on an understanding of the people, their practices, their capabilities, their history, and the work needing to be done. The design has to be model-based, even if you model the work system in a text editor. That's the way we started our Brahms models, by the way, as an outline of the people and their roles, the software systems they used, the objects in their environment, and the places where they work. That's the basic vocabulary in the Brahms framework.

In the ideal form, there's a simulation of the work process. You can run scenarios that you anticipate will arise and can run a range of others that are outside normal bounds. You can make predictions. You can discover vulnerabilities relating to workload gaps, timing issues, overall productivity, costs, and so on. We instrumented Brahms to be able to generate statistics like these to answer questions that managers and engineers might have about the future work system. A manager might want to know whether changing people's roles and technology together would be effective for the anticipated workload, for example, or whether productivity could be doubled in practice using the new technology.

In the scientific design approach, you develop prototypes, put them out into the world, and experiment. My ideal, which I was never able to get the funding to do, is to simulate the scenarios that you plan to test in the field with prototypes, and then see if the performance of the overall people-technology-environment system matches predictions from the model. Your experiments give you the information you need to refine the model.

The other point that I want to emphasize is about the consequences of not getting this right, especially for very complex projects. I mentioned self-driving cars and aircraft automation. There's also the example of the Navy's Littoral Combat Ship, which intended to use new automation technologies to minimize the size of the crew. It failed to meet expectations, and I think that a failure to appreciate the relationship between the people and the technology is the reason for that. Billions of dollars were invested in ships without a work practice simulation and without incrementally prototyping the new work system concepts.

What these R&D failures had in common was the use of idealized models that didn't properly account for the actual operational context. The Navy didn't account for the range of scenarios in which systems and people would need to interact. For example, what happens when the crew is diagnosing and repairing advanced automation while at the same time they are responsible for sustaining other operations? Given the complexity of today's automated systems and the tendency to automate everything people do, just for the sake of it, this is becoming a critical engineering problem. The vulnerabilities, risks, and costs are rising together.

JE: Can you talk about how this can be done in practice?

BC: In our original work on Brahms with NYNEX S&T, we had the idea that the model would be a conversation piece. It would bring together people from different jobs and different disciplines to understand what might happen. The model objectifies the work, so the discussion is not about individuals and their behavior. It is something we can point

to, which facilitates reflecting and learning about the perspectives of others. If that is all it does, that's great.

But a work practice simulation can do much more. It can help to answer questions that the designers and managers have about the new work system. It can provide quantitative predictions that are useful in the design process. It enables us to verify (or invalidate) design hypotheses about personnel, resources, and throughput.

The power of dealing with emergent phenomena was especially evident in our NASA Ames simulation of the 2002 mid-air collision over Überlingen. In the air traffic situation, planes are moving, air traffic controllers are moving between workstations, a radar system that shows objects moving on the screen may represent what was happening 10 seconds ago. These space-time interactions are extremely difficult to predict on pen and paper. What happens if a third pilot repeatedly distracts the controller, for example, and the short-term warning system is disabled because of maintenance in the control center?

A simulation that captures the emergent events excels for problems like this. You can crank out hundreds of scenarios, varying the timings a bit, and you can estimate how often planes are going to come within some distance of one another. You can see how often an air traffic controller is going to be looking in one place and not seeing something that needs to be responded to. You can start understanding vulnerabilities in the work system such as cognitive complexity, inflexible protocols, lack of cross-training, failure of supervision, and so on. An activity-based, distributed agent simulation helps to understand the work at that more detailed level. But you have to include the people in the design model! The Traffic Collision Avoidance System (TCAS) was validated without considering the complications when TCAS and the air traffic controller give pilots contrary instructions.

JE: You've done this kind of simulation in a lot of settings, including some extreme environments. If you were considering designing work in a more routine world—maybe work in a call center or the work of a field technician—what would you do?

BC: I would always start with observations of the work—preferably by an anthropologist working with a computer scientist who happens to have ethnographic skills. Based on the observations, the two of them would begin the modelling, not at a code level, but at the level of agents, objects, groups, locations, communications, and the different activities taking place—things like moving, perceiving, and detecting. It's important to start with a diagram laying out the places and an outline of the people, objects and automated systems that matter. Then you need to map out their activities.

JE: You've been telling this story for a long time, but the world of work design is changing very slowly. At the same time, the world of AI and automation is changing very quickly. What do you think the future will bring?

BC: Well, eventually failures in the real world are going to force a change in the way we do design. One reason for hope is the initiative at the National Science Foundation (NSF) to fund virtual AI Institutes. NSF is going to fund 8–10 consortia in the current round, and I've worked with one of them. My encouragement lies in the fact that the people on the team are concerned about the same issues that we've been discussing. They start with the ideas that AI should provide tools to amplify people's capabilities. They understand the importance of distributed cognition. And they are addressing important social problems, like emergency response, health care, and more broadly, the ethical use of AI.

JE: Wow! That's wonderful. I need to learn more about that.

BC: Progress takes time. There is a lot to assimilate with a new, comprehensive design approach. To do this work, you need to understand people and technology together as a system. I expect that there will be new approaches to sociotechnical design that draw on what we have learned from neuroscience, computer science, psychology, philosophy, and the social sciences. It is slow, but we're getting there. Engineers are more ready to talk about work practices and what used to be called the "soft stuff." The d.school at Stanford is one of the places leading in this area. But a systems thinking approach began a long time ago, in the 1950s, when Buckminster Fuller advocated a "comprehensive anticipatory design science."

There's another development that I find promising. About a year ago, Yolanda Gil and Bart Selman wrote *A 20-Year Community Roadmap for AI*. It is very good. Many researchers are beginning to focus more on ethics nowadays, and the roadmap reflects that. There is a developing understanding that we cannot just deploy AI programs without concern for biases in the model data—for example, a flaw that recently surfaced in some face-recognition systems.

JE: A lot of the work on AI ethics is still focused on issues like bias and manipulation, though, not on work design.

BC: It's starting there, but I hope that it doesn't get stuck there. The larger problem is how people will use these tools, and that's what work system design is all about. In addition to model validation, design for responsible action by people needs to be part of the design of AI tools.

JE: Yes. I would like to see a focus on how to create "good work"—not just work that doesn't exploit, but work that engages the intelligence and commitment of people.

BC: Here's my current hypothesis: the necessary change in design practices is going to take time. Paradigm shifts always do. Success will depend in part on presenting ideas relating to work, people, and technology early in the education of technologists. Engineers need to be trained to use the new scientific design methods because the risk analysis and

simulation methods they're using now for megaprojects aren't working. It's likely that they're not working for other projects either, by the way, but when they fail, it's just not as visible. These scientific design methods are going to involve observation; they're going to require models and "total system" simulations; they're going to be empirical, refined over time, developed incrementally, and designed using participatory approaches. We know the methodology inside out; some aspects of it have been used in corporate research settings for 30 years. Now we need to spread the word.

But we can't just throw tools at the engineers. We have to teach them to be reflective about the ways that they currently solve these kinds of problems and what the limitations of their current practices are. They need to understand the things that are now invisible to them in order to be motivated toward doing better design, and they need to understand specifically when and why existing modeling tools are failing. A big part of my work has been to make the invisible visible in models used for the design of technology and to instill a scientific way of thinking about these things.

R&D MANAGERS AND INNOVATION LEADERS NEED A **LEVEL OF JOB PERFORMANCE** *THAT IS BETTER THAN SATISFACTORY.*



The candidates at Innovation Research Interchange deliver consistent excellence – a standard which can only be met with continuous access to state-of-the-art skills and continuing education. By leveraging the power of a trusted association, you tap into a talent pool of candidates with the training and education needed for long-term success.

Don't miss this unique opportunity to be seen by an exclusive audience of the best and brightest in Research & Development and Innovation.

Visit the IRI Career Center to post your job today!

CAREERS.IRIWEB.ORG

Visit the IRI bookstore!

Access the best IRI has to offer in electronic or print-on-demand formats. Browse *RTM* reprint collections, white papers, and special offerings, all selected and compiled to offer a range of perspectives on central issues in the management of technological innovation.

http://www.iriweb.org/bookstore