

Snow's Two Cultures Revisited:

Perspective from Design in Human-Computer Interfaces

Bill Buxton

University of Toronto and Xerox Palo Alto Research Center (PARC)

In 1959, Sir Charles P. Snow presented a landmark lecture on the relationship between the science community and that of the arts and humanities, or "literary intellectuals." In his essay, Snow characterized these two communities as having lost even the pretense of a common culture. That is, they had lost the ability to communicate on any plane of serious intellectual endeavor. This he argued from the perspective of creativity, "intellectual life," and normal day-to-day living. This was the essay, then, that coined the term "the two cultures," to characterize the polarization of these two communities.

In what follows, it is my intent to take a renewed look at the relationship between these two cultures. The argument that I am going to make is that the technology of computers and telecommunication has brought about a potential meeting ground for these otherwise polarized worlds. This can be stated even more strongly. Consider the interface where humans meet technology. My belief is that the problems that emerge are so critical to society at large that they require and compel both "cultures" to work together in order to resolve them in a socially acceptable way.

My intent, however, is not to present a detailed map of this "meeting place." Nor is it to present some reductionist treatise on the dangers to society of technology and how they can be solved. Rather, my plan is to jump (somewhat cavalierly) among anecdote, musing and argument to identify some of the landmarks that make up this "place." My hope is that the rough sketch that emerges will provide some sense of how the roles of the two cultures can converge and complement one another. And since I am speaking primarily to those from the arts, my emphasis will be mainly on the role that they have to play in this meeting.

Before proceeding, let me make the aside that gulf articulated by Snow is in ample evidence today, even (especially?) in the very area where I have claimed the potential meeting ground exists. We see it expressed as the old "them and us" attitude, as manifest in complaints such, "*Those engineers! Those programmers! We're the intellectuals; this is how we write/paint/draw. Those cybernerds have no concept of what we do, much less know how to build our tools.*" I would simply say that the "them-and-us" attitude of which this is a caricature is destructive, unnecessary and extremely counter-productive.

But if there are two cultures, and we think that the gap should be bridged, then why can't we just do so by education? After all, isn't that one of the goals of a liberal arts degree - to give us the opportunity to balance our literacy in the arts with that in the sciences? Would that it were so simple. Snow's original lecture was largely directed at the issue of education. But in it he also pointed out why this was no magic bullet. Simply put, the modern world is just too complex. Even by the eighteenth century, when the Encyclopedists were active in France, the complexity of the world was beyond any individual to absorb. As Snow pointed out, trying to use education to create a new generation of neo-Renaissance men, or women, is not a viable approach to the problem.

Yet, parenthetically, that is close to one of the "lies" being sold today with multimedia; namely, that you, too, can be a Renaissance man or woman. You can be a good graphics designer, sound designer, scriptwriter, salesperson and scientist, all at once. Of course this is nonsense --unless you're far brighter than me or anyone else that I know. Renaissance teams? Yes. Individuals? No.

Let's look at such teams then, using my own initiation as an example. Back in 1969 I saw a computer for the first time, and used it to create a film soundtrack. This computer was really a music workstation: it had four voices real-time multi-timbral digital sound, used common music notation, had color encoding of the parts, and let me use both hands to efficiently enter music (using a chording keyboard on the left hand for note durations, and a pair of thumb-wheels on the right hand to specify pitch and entry point in time). Even the tape recorder was under computer control.

I worked on that system at Ottawa's National Research Council. On the same computer was also one of the first interactive animation systems, which was capable of performing automatic in-betweenings from key frame to key frame. It was the system on which Pierre Foldez created his award winning animation *Hunger*. Pierre did his animation on the system during the days, while I worked on my film score at night.

I spent a week with this system, finished my score, and went home. I thought that was what computers were, and was impressed: walk in, do your score and walk out. What could be easier? I was sold. So although I was a musician, I decided to take a computer course. That's when I hit the wall and slid into the culture gap!

I had the dubious pleasure of studying "Burroughs Extended Algol." But this had nothing to do with my experience in doing the sound track. I couldn't figure out how to reconcile these two experiences. Nevertheless, I persevered because my initial experience was so strong and positive that I *knew* that there was another way (I had experienced it). In retrospect, this experience laid the basis for my subsequent life in design and research: I was just trying to get back to where I started.¹

I was lucky. My first experience occurred with people who were extremely enlightened about human-computer interaction and sensitive to the arts. They also had the wit to realize that the creative arts offered an excellent base to find guinea pigs for studies in the design of human-computer interfaces. A lesson that I learned from them and have benefited greatly from over the years was, who else does creative stuff and is foolish enough to stay up all night and work on a computer? Not your typical dentist or lawyer! An artist understands how a pencil should work and feel. Likewise, musicians make a lifetime pursuit out of expressing powerful ideas through a mechanical intermediaries. They understand what an "instrument" is. In short, to design an interface that has the qualities of a good sable brush or a quality musical instrument is something to strive for, regardless of application domain.

I use the analogy with brushes and instruments because they are the intermediaries between what is in the mind and its realization in the external world. Like computers, these technologies also have "users," *but it is almost an insult to describe them as such*. Why? Because of the skill that they bring to the task. Where these tools differ from most computers, however, is in the quality, finesse, and subtlety with which they can capture this skill. And there lies both the core of the problem and the key to its solution.

Consider this: I would wager that the first violinist of the Los Angeles Philharmonic uses a violin bow -- I'm talking about the bow, not the violin -- that costs more than a Sun workstation. Think about that, in terms of the message it conveys regarding the relative importance accorded to quality tools (especially given the relative income of artists and their ability to afford them). Yet expensive computers that perform "serious" functions don't let you do what you can do with a 5¢ pencil, such as draw a line whose thickness depends on how hard you push. Even on a so-called "graphics workstation" you have to buy special hardware and software to do a dumb little thing like that.

But so far, we have been talking about the problems of artists as users of computers: sort of a "consumer as victim" type of whimpering. That's not my intent. Where we are going is to explore territory where aspects of the arts are applied to solve some of the problems that we have been confronted with. Along the way, my hope is to make the point that these solutions benefit the community in general, not just those from the arts. The bridge is somewhat that of the art of the instrument builder, or "luthier."

¹ For more on the NRC system, and its impact on my work, see: Buxton, W. (2008). My Vision Isn't My Vision: Making a Career Out of Getting Back to Where I Started. In Thomas Erickson & David McDonald (Eds.). *HCI Remixed: Reflections on Works That Have Influenced the HCI Community*. Cambridge, MA: MIT Press, 7 - 12. <http://billbuxton.com/MyVision.pdf>

Let me take a hot topic for my example: *Scientific Visualization*. Today we're in a computational and communications revolution, which often is called the "information revolution," and which is frequently compared to the Industrial Revolution. However, this "information revolution" is a myth! Instead, what we really have is a *data explosion*. But data does not become "information" until it informs or can serve as the basis for decision-making.

Over the last three years, every magazine that covers computer graphics has published articles about scientific visualization. Scientific visualization involves graphically rendering complex data in such a way as to make pertinent parts and relationships within the data more salient to the viewer. The idea is to tailor the visual presentation to take better advantage of the human's ability to recognize patterns and see structures. From this perspective, scientific visualization is one step in moving from a *data* to an *information* technology. In pursuing the thread of our discourse, the good thing about this is that we can see a clear role here for those in the visual arts and graphics design, since it is precisely the skill set that they possess that forms the essence of the technique. The bad news is that scientific visualization doesn't go far enough.

To begin with, why just "scientific" visualization? What we should be pushing for is for similar techniques to be applied in other domains than science. But even if we restrict ourselves to the term "visualization," we still have problems. The visual system is not the only sensory modality that can be used to extract meaning from complex data. For example, without the use of the audio channel, one would be seriously handicapped in performing complex tasks in the everyday world, such as driving or crossing a busy street. Likewise, I believe that we are seriously disabled in terms of our ability to function in the domain of information systems due to the impoverished use of sound. So let us pick a term that captures the full sensory richness of the user. Let us pursue *perceptualization*, rather than visualization. And in so doing, we now open up the possibility of building upon the skills of musicians and sound designers as well as visual artists.

Finally, before leaving the subject, there is one last problem. Visualization and perceptualization conjure up too much the image of the human as mutisensory sponge: information is presented to be absorbed. The problem with this is that this is not how humans learn nor how they develop an understanding of complex phenomena. The process is rather far more akin to exploration. We learn by becoming actively engaged with our material. So having started with scientific visualization, my argument is that what we should be striving for in *interactive* perceptualization.

When we think along these lines, we see less tension between art and science. We start to expose a meeting ground for the two cultures. The ability to tailor systems to reflect how we sense, think and problem solve requires contributions from both the arts and the sciences. Neither has a monopoly. Each needs the other.

The principle design sensibility underlying the ideas just discussed is one of designing to exploit the existing skills of the user -- the skills acquired through a life-time of living in the everyday world. If we embrace this skills perspective, we open up the door to understanding the ubiquitous (but misdirected) objectives of "ease of use" and "user friendliness." As objectives for design, both of these are poorly formulated. Our focus on skill leads us to a much better formulation of our design objective: *to design systems that accelerate the process whereby novices begin to perform like experts.*

Expert performance isn't simply faster; it's not just a quantitative difference; it's a *qualitative* difference. The way we structure our activities - how we think, how we formulate problems, how we structure our activities, how we move our hands and so on - differs between expert and novice. When we formulate the objective in this way, whether designing educational software, drawing programs or music systems, we are forced to ask, "What are the essential skills involved in what I'm trying to support or encourage?"

By examining the differences in how an expert works as opposed to how a novice works, we can better design systems that channel behavior along the desired path. We can start building strong, specific systems tailored to support specific functions.

But there are different types of skills and expertise. How do we take the general dogma expressed above and establish some design criteria that reflect these different skill levels? My approach is best described using a metaphor of three mirrors (Buxton, 1991). In this metaphor, we think of the technology as being three mirrors, each corresponding to a different type of skill. From this, the quality of a design is directly related to the fidelity with which that skill set is reflected.

The first "mirror" reflects the user's motor-sensory system. How well does the system reflect the sensory and motor skills of the user? Typically, the image is rather distorted. If you held up a mirror to a typical microcomputer, you would have to assume that the user either has one eye with monochrome vision (or sees only fully saturated colors), has ears with a frequency response of about 200 Hz, understands no spoken language, has no legs, and has one arm ending in about 60 uniform-length fingers that only have two joints. Given the evidence, these conclusions are as reasonable as they are wrong. Either we should change our bodies or change our machines.

Blindfold a few good flautists and hand them two flutes: a poor one and a good one. Without playing a note, they know how much respect for their skills is built into the tool, simply by touching it. The distortion at this level reflects a lack of respect for the skills of the user - distortions that would never be tolerated from an instrument designer.

Let us now look at the second mirror. In this mirror we ask, "How well does the computer system reflect our cognitive structures: how we think, how we learn, and how we solve problems?" Unlike things at the motor sensory level, these structures are invisible and harder to quantify. Yet there is ample evidence in the frequency of errors, and difficulties in solving problems that the distortions at this level are as great, or greater, than those reflected in mirror one.

The third mirror is among the most important, yet perhaps the most neglected. This is the mirror that reflects our social skills. With this mirror we ask, "How well does the design reflect how I go about my day-to-day work and how I interact with people and groups?" Virtually all computers were designed for use in a one person and one computer face-to-face configuration. But that's not how life is. Nor does it reflect how we work. In our daily life, we meet other people. We have social interactions. We jump from idea to idea, from person to person, and group to group. The degree to which computer systems support these types of interaction is impoverished, to say the least. This third mirror might as well be painted mat black, for all the information that it reflects.

But it doesn't have to be this way. Again, we can go to the arts for an example. Think of a musical instrument that is designed to be played by four people and that resembles a non-competitive game (because the only way you win is not by competing with the other three people but by collaborating with them). Such an instrument exists. It's called a string quartet.

The string quartet is not just four instruments. When performed properly, it represents one instrument played by four people. Stated differently, it represents one technology capable of capturing the collective efforts and skills of a group in order to solve a particular (musical) problem. What we need, and is starting to emerge from the world of "groupware" and "computer supported collaborative work (CSCW)," are technologies that can capture such group efforts on a broader class of problems, and do so in a manner that reflects and exploits existing human social skills.

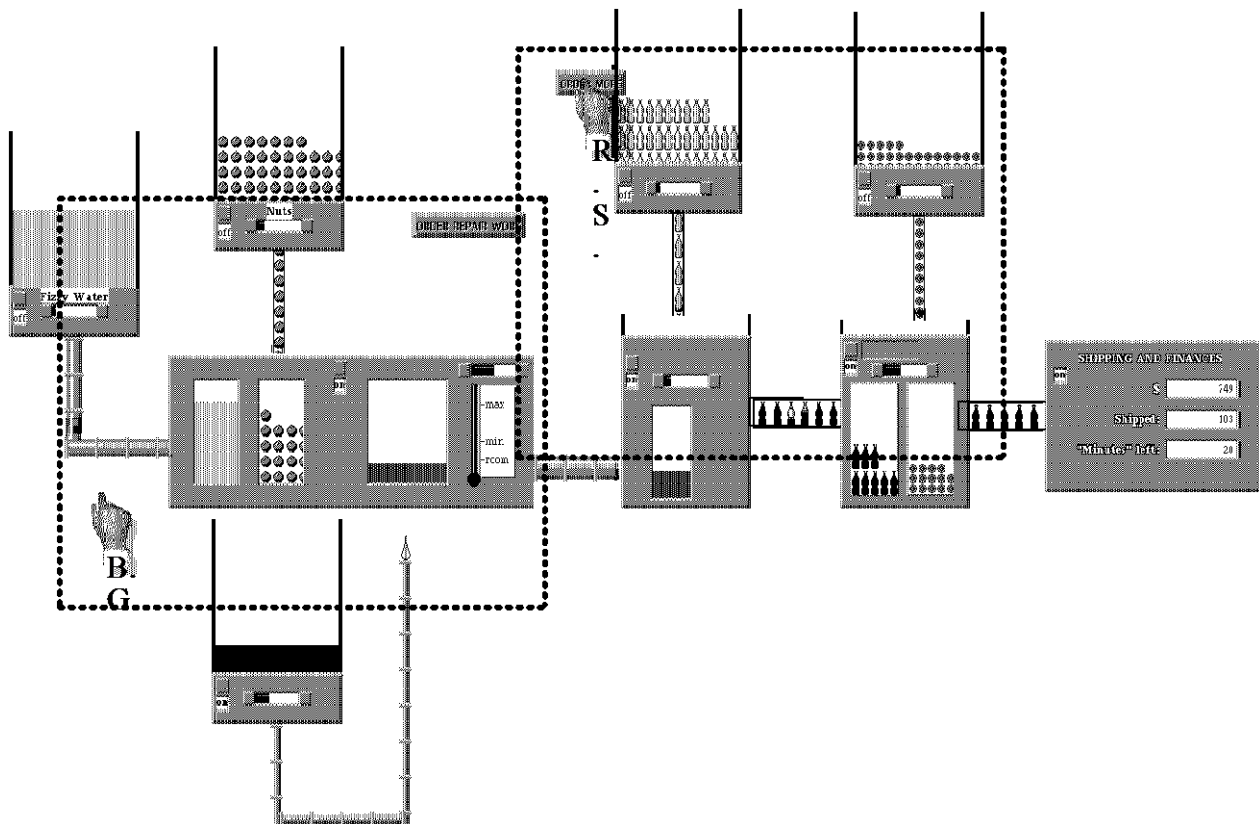
This discussion of group activity brings us full circle back to Snow and his observations that our hopes should not lie in education breeding a new generation Renaissance person. If we accept this notion, then our only option lies in teamwork and collaboration, and in technologies that can help support, foster and encourage such activity.

We must remove the distortion from the three mirrors. This task requires expertise from both of the "two cultures." Therefore, we need to take steps in terms of design and our organization to build up the too long neglected trust relations between the two groups. Let me finish, then, with an extended example that epitomizes the creativity and benefits that can result when we start to let this happen.

The example that I have chosen was undertaken at Rank Xerox's Cambridge EuroPARC lab. This is a satellite of the famous Xerox Palo Alto Research Center (PARC). It was carried out by two colleagues,

Bill Gaver and Randall Smith. The goal underlying the work was to study collaborative computer mediated work and problem solving by two individuals in separate locations. Subjects were connected by a reciprocal audio/video link, so they could both see and speak to each other. In addition, the task that they were jointly working on appeared on the screens of their computers, which were connected by a local area network.

The task that the subjects had to perform was to operate a simulated cola bottling plant. (Since the underlying system was called the Alternative Reality Kit, or ARK, the study was known as ARKola). A key aspect of the study was to investigate the effectiveness of non-speech audio as a means to augment the subjects' ability to operate the simulation efficiently. Some subjects ran the simulation with, and others without the non-speech sonification.



The figure is a screen snap of the simulated bottling plant used in the ARKola experiment. Each of the shaded rectangular objects represents one of the machines in the factory. For example, the left-most one is the source of "fizzy water," and the right-most one is the "shipping and finances department." The two larger dotted rectangles illustrate what part of the factory each of the two "operators" is currently viewing. Operators can move their view independently, but neither can see the whole factory at one time. At the moment, there is little overlap between the two views. The two hands seen in the figure (lower left and upper center) are controlled by the operators' mice. Which is controlled by which operator is indicated by the initials under the hands ("B.G." and "R.S." respectively). Given that the factory was larger than the operators could satisfactorily monitor visually, the sounds made by the various machines played an important role in helping the operators detect, diagnose and repair problems. (Figure compliments of Bill Gaver, Rank Xerox EuroPARC.)

According to Bill Gaver,

In this experiment, we were interested in exploring the usefulness of sound in a large-scale collaborative environment. In particular, we were interested in whether sound would be useful in

helping subjects with a number of simultaneous, interdependent processes, some of which were not visible on their screens.

We implemented the project in the Shared ARK environment, a large-scale collaborative system developed by Randall Smith of Xerox PARC. The task we gave subjects in this environment was to run a simulated cola bottling plant. Their goal was to work with their partner to make money and produce as much cola as efficiently as they could.

The cola bottling plant consisted of nine machines connected by a series of pipes. Basically, water and nuts were pumped from tanks into a machine called the "cooker." The cooker heated the cola; when it reached the right temperature, nuts and water were combined to produce cola, which then was added to a holding tank. At the same time, some of the cooked cola was pumped out of the cooker and into the bottle-filler machine. This machine filled bottles with cola and sent the full bottles down the line to the capper machine. The capper sent capped, full bottles to the end of the factory line, to a "shipping & finances unit" machine.

The shipping and finances machine counted the successfully produced bottles and paid the operator a set sum for each bottle. It also took away money when new supplies were added to the holding tanks. The current balance appeared on the computer display. Again, half the machines make the cola, and the other half bottle and cap it. Bottling supplies run out occasionally and must be refilled by the operators. Refilling tanks cost money. The operators' goal was to produce as many bottles of cola as possible without wasting too many supplies.

The bottling plant appeared to operators on the computer screen and display. Not all of the plant was visible on screen at one time. Instead, operators used their hands to scroll over the virtual underlying surface of the "world," and collaborating partners saw different views of the world and could divide tasks so they oversaw different machines.

Sounds helped the operators achieve their goals. Some sounds provided information about the status of on-going processing. Other sounds indicate trouble, such as empty bottles reaching the shipping & finances unit. For example, the heater made a sound that changed in pitch as the temperature rose. The bottle-filler machine emitted a dripping sound to indicate a tank overflow--which meant a waste of material, and a waste of money. When the machines ran smoothly, another sound indicated successful operation.

Can non-speech audio cues actually help operators oversee their work tasks and collaborate better with other people who see a different part of their working environment? I believe so. ARKola represents a "multimedia" system where the use of our senses and interactive media combine to help accomplish fundamental tasks, instead of serving as some secondary artifact, or lollipop.

The design of ARKola system required a great amount of musical skill in designing the sounds. Although it's not "Art," nevertheless, it represents a meeting ground between the two cultures which can result in improved design.

In 1959, Snow used knowing the second law of thermodynamics as an example of basic scientific literacy, equivalent to having read Shakespeare as elementary literacy in the humanities. While even today only a minute fraction of literary intellectuals would be able to speak at all about thermodynamics, the "two cultures" of today have a technology-based common meeting ground and vocabulary. In addition, given the potential of emerging technologies to impact our society, I believe that taking advantage of this meeting ground is compelling and essential for the health of the society.

Certainly, as someone from the arts working in technology, one is too easily cast into the role of being a champion or advocate of technology. But given the ceaseless advance of technology, one can either attempt to achieve a degree of literacy and help shape its evolution in humanistic directions, or bury one's head in the ground. With an equally healthy dose of skepticism and optimism, we arrive at a healthy degree of *skepticism* in approaching the topic: an approach based on respect for the skills and contributions from both cultures.

REFERENCES

- Buxton, W. (1991). The three mirrors of interaction: a holistic approach to user interfaces. *Proceedings of Friend21 '91 International Symposium on Next Generation Human Interface*, Tokyo, Japan, Nov. 25-27, 1991. Also appears in English and Japanese in: Buxton, W. (1992). The three mirrors of interaction: a holistic approach to user interfaces. *Industrial DESIGN 157*, Japan Industrial Designer's Association, 6 - 11.
- Gaver, W., Smith, R. & O'Shea, T. (1991). Effective sounds in complex systems: The ARKola simulation. *Proceedings of CHI'91*, 85-90.
- Pulfer, K. (1971). Man-machine interaction in creative applications, *International Journal of Man-Machine Studies*, 3, 1-11.
- Snow, C.P. (1964). *The two cultures: and a second look*. Cambridge: Cambridge University Press.