THE FUTURE MERGING OF SCIENCE, ART, AND PSYCHOLOGY

Marvin Minsky

1. Reason and Intuition

One day, walking through the laboratory, I met a colleague and asked about the course he was teaching. "I got into some trouble today," the professor said. "What I was explaining seemed so obvious that I couldn't see why the students couldn't understand it. I've taught this subject so many times that I can't remember which parts are hard." What do people mean when they say that certain conclusions are logical while others are intuitive? Are these really different types of thought? It seems to me that the differences lie less in the styles of thought themselves than in the degree to which we understand how they work. For, in general we're least aware of what our minds do best. We can usually explain, in much detail, how we perform a mathematical calculation or procedure; this is because we already know how to describe it in terms of its simpler parts and relationships.

But we're speechless when asked to explain how we speak, or of how we hear or how we see. Clearly, these accomplishments of everyday life involve immensely intricate brain-machines — yet to us they proceed with no effort at all. What makes such complex processes seem so simple to us?

To the computer programmer, one possible answer seems obvious: that our most practiced and most "common-sense" skills must be just the ones that we have converted or precompiled into procedures and scripts that can be executed unconsciously. It is only when those systems fail that we start to engage the special types of procedures and memory systems that comprise what we call "consciousness." Thus, we're more aware of processes that don't work well than of those that tend to work flawlessly. This means that we cannot trust our offhand judgments about which things we do are simple, and which engage complex machinery. This must apply equally to artists and scientists. All experts accumulate towers of skills, building new ones on top of older ones. At every stage we become novices again, finding new goals (and hence, new obstacles) and then seeking ways to deal with them. But as we develop each layer of skill, we revise and refine the layers below, to make them work more efficiently, more quickly and automatically, and with fewer intrusions on consciousness. Thus as we grow in mastery, our highest level processes grow further and further away from the earliest machinery — and it becomes increasingly more difficult for those higher level systems to detect, discover, or describe how their lower level systems work.

The more we improve our mental skills, the less we know about how they work. A novice must deliberately consider every step and consciously keep track of what goes on, perhaps by using elaborate tree-structured representations. But the process of becoming an expert includes learning how to replace those clumsy but thoughtful procedures with more smoothly working script-like structures, and to substitute habit-like, direct connections for those elaborate representations based on complicated directories, pointers and cross-references by much more. Then the new, more streamlined procedures rely less on thoughtful (and painful) analysis and more on automatic processes. This leads to the seeming paradox: a novice may not perform very well, but can better explain what was done, while a master seems better to "know what to do" but knows less about how it is actually done.

Then analysis seems superfluous, and the technique that works seems "obvious." I often hear such performances attributed to intuition — that is, to aspects of mind that cannot be explained. And indeed, no matter that this idea may be basically wrong, it helps us to restrain ourselves from trying to understand ourselves. We usually work most efficiently when not distracted by self-analysis, which tends to disrupt all those craft-bearing scripts. What happens when you attempt to observe your own mind at work? At the very least, this must divert some of the precious short-term memory that otherwise could be used for the task at hand. And it should be even worse if you succeed at successfully probing into those towers of fine tuned processes — and transiently replacing them by clumsy simulations. This surely should render you awkward and uncomfortable — and make you fear loss of your mastery.

But I suspect that thinkers must suffer in order to grow, and those quick-working scripts can restrict that growth, for they tend to become inflexible. Each unthinking skill, eventually, may close off possibilities. To escape that narrowing of range, one must learn to enjoy the pain of retreat and force oneself to experiment. Each novel way to wield brush or how must make you feel ungainly at first — and force you to suffer for hours and days. But eventually the new techniques will refine and compile — until once again you can't tell what is "hard," and performance then becomes painless again.

My book *Mentopolis: The Society of Mind* takes a more contrapuntal view of pleasure, pain, and suffering, in which the pain of growth is seen as largely an illusion because, at the same time, other parts of the mind may silently rejoice at the experience of learning and development. On the other side, the sense that an intense pleasure is generally satisfying may be equally illusory, when it comes from some smaller fragment of the mind having seized control, depriving all the other mind-parts of opportunities to achieve their goals.

2. Goal and Description

What do artists try to do? Some seek to influence other minds by communicating views or ideas, or inducing moods or attitudes. Other artists are less concerned with specific impositions, and more with trying to help others to free themselves from old habits and presumptions, so that they can find new concepts within themselves, or novel ways of seeing things. In any case, the artist wants to have some effect on the viewer's mind. Now, sometimes the artist knows just what to do, and starts with some more or less well-formed sketch. But, usually, when you compose, your own objectives are growing, too; perhaps you begin with little more than some cluster of conditions and requirements. Then your situation might be better described in terms of a search than an image, because your goal is to find or construct something that satisfies those requirements. From this point of view, the artist's pursuit is no different from that of other seekers: one must propose tentative solutions, examine them for virtues and deficiencies, and proceed to change, adjust, and refine them in accord with one's current goals.

But is it proper to think of an artist as a problem solver with goals? Certainly, some artists dislike to describe themselves that way. But this is not good evidence, because artists are people — and people, in general, find it hard to acknowledge (or even recognize) their deepest goals. As Freud observed, our most basic motives are often the ones of which we are least aware. Furthermore — and Freud did not emphasize this — to express or describe what happen inside the mind is itself an extremely difficult problem; indeed, great depicters of mental activity are as rare as great composers, painters, or architects.

No matter what one's purposes, perhaps the most powerful methods of human thought are those that help us find new kinds of representations. Why is this so important? Because each new representation suggests a new way of understanding — and if you understand something only one way, then you scarcely understand it at all. Perhaps this is why the arts so often precede the flowerings of culture. For what, indeed, is Art itself but the Science of discovering new ways to represent things?

We can see this operating in the difference between producing a particular work of art, and developing what we call a distinctive and coherent style. This happens when an artist invents, not merely a new description of some particular thing, but a new kind of representation — a new type of viewpoint or way to describe. And once this happens, the artist's experience undergoes change: the audience may see the performance as art, perhaps intricate and mysterious. But the artist soon comes to regard it, instead, as an obvious, natural way to proceed. It is something that he now knows how to do — though he may be unable to say it in words.

Then what distinguishes an artist from any other kind of scientist, craftsman, or specialist? Sometimes, surely, nothing at all. And, sometimes, simply, a lack of constraint. (An engineer's engine has to run, but not an artist's depiction.) But, in my view, what characterizes many artists is quite the opposite: they seem impelled to accomplish, simultaneously, impossibly many incongruous goals!

I recall an engine in one of Diego Rivera's murals that must have been intended to affect many sides of the viewer's mind: to convey a constellation of statements and attitudes about the workers (and about their attitudes) concerning the engine, the ambient technology, the factory environment, and their relations with other inhabitants of that institution. So many things to say at once. Thus, while a scientist or engineer usually confronts a certain fixed set of clearly defined requirements, an artist entertains more goals at once than ever can be satisfied — and therefore also must face yet another kind of problem, too: which goals to pursue or abandon?

3. Scientific Understanding of Expressive Art

In ancient times, the visual arts were confined to rather static modes, such as painting and sculpture. But in the past century, the art of Animation emerged and has steadily grown in popularity. If not for its enormous cost, it might have become the dominant form in many domains such as entertainment, education, and technical exposition. In recent years we have gained access to powerful yet inexpensive graphic engines — yet animation is still too expensive for casual use, because programming it is still very hard. And this has produced a typically modern predicament. As soon as they sensed the promise of being released from the limitations of the old technology, the artistic community — and particularly, animators and musicians — faced bondage to a new and demanding god. They felt obliged to make themselves more technical, in order to invent the science they would need to foster their art.

In older times, a painter had only to cover a single canvas. But to make a five-minute film, an animator must produce nine thousand frames. How to paint so many scenes? The work is reduced if one can reuse what already appeared in earlier frames; this can be done with technology that enables the artist to work in terms of differences — by using interpolation, "in-betweening," mathematical interpolation functions, and other ways to generate scenes. But still there are many new problems to solve. How to make a person walk — that is, saunter or stroll or strut or stride? How to make a person throw — or toss or hurl or heave or fling? It seems incredible how well some artist-painters of those older static forms succeeded in expressing such things. What are the tricks those artists used to produce such compelling illusions — and what new tricks could animators use?

In the past generation, many workers have devoted themselves to these matters — of how to adapt computers to the purposes of artistic expression. We ought to acknowledge the sacrifice made by a generation of potential artists in order to accomplish this! How often so many of us complained, in those pioneering decades that "these new electronic synthesizers sound too mechanical" or that "computer art is elegant but lifeless." Of course there would be deficiencies in the early forms of such radically new technologies and tools, and their languages and idioms. We ought to thank, not ridicule, the generosity with which so many potential masters — both in music and in art — unselfishly renounced the joys of artistic composition to endure the pains of evolving our new instruments and programming techniques. Of course, we could regard the instruments themselves as instances of new forms of art.

But perhaps we are approaching new ways to understand art itself! Consider that, in earlier times, the working of the human mind has been an almost total mystery. The science of Psychology itself is only now a century old, and Computer Science only half of that. Surely, in the next hundred years, the twin enterprises of Cognitive Psychology and Artificial lntelligence will reveal a great deal more about how minds work. Thus, in the past, we had no choice but merely to appreciate and celebrate when certain individuals mastered new secret skills of expression — no matter that they could not explain them. But, it seems to me, this must change, as the world increasingly accepts the products of the works of art whose production is part or wholly based on computational processes. For then, the audiences will feel increasingly concerned with such questions as "What mechanisms or technical procedure produced that expression?" and "Why did that expression have such an effect upon me, my minds, and my emotion?" and, finally, "What kind of mechanism or procedures must exist within myself, in order that effect or experience could be produced?"

What I mean by this is that the "mechanization" of artistic production must eventually and relentlessly lead us to become involved with these new kinds of personal concerns. On the creative or productive side our artists, in order to program their new machines, will be virtually compelled to formulate, in more conscious detail than ever before, what it is that they want to produce.

And equally, on the receptive side, the audience will be forced to confront those questions about what makes them react. Thus, from both sides, this brings us back to the initial confrontation of this essay — namely, that paradoxical recognition that the things that seems to us most "obvious" are often the hardest ones to explain. We have needed to have that same experience, from the earliest days of research in Artificial lntelligence. Some of very first such computer programs were able to solve problems that people considered to be very hard — yet even today our programs are unable to manage most "commonsense" human activities.

Why was it easier, historically, to construct seemingly "expert" computer programs than to make progress in domain of "common sense?" One reason is that those technical domains frequently are actually simpler in nature. For example, a fundamental domain of mathematics or logic may be based on only a handful of assumptions and rules of inference — whereas the art of cleaning a house requires many kinds of knowledge about a great many different kinds of objects and situations. Another reason is that many "technical" subjects have already evolved powerful, specialized artificial languages and notations — whereas we have few such well-developed languages for activities that common people do! This is perfectly natural since no person ever needs to precisely describe the things that everyone learns in childhood, since everyone else already knows them already. (Everyone knows how to walk, or to laugh, or to distinguish the differences between cats and dogs — so no one ever needs to be told how to do it.)

But we cannot exploit any such common experience when we want our machines to help us express such things — for example, when we want to use a computer to assist us with making animated pictures. This is because the initially empty random access memories of our computers know nothing about commonsense human affairs — about people's emotions, motives, goals, frustrations, and joys; or about the systems that control the motions and actions of the human body.

What would be the ideal technology for an ambitious animator? Imagine being able to present your machine with a screenplay and a storyboard — and have it generate the film. To do this, the filmmaker wants to control not only what the animated actors do and say, but also the intended effect on the viewer's mind — that is, to evoke a certain sequence of ideas, attitudes, feelings and emotions. Now we certainly cannot expect software to help with this, unless it has access to knowledge about rendering physical forms, motions, and actions. But in order to induce the desired subjective reactions, our machines must also be equipped with descriptions and models of human feelings, goals, and intentions, or those simulated actors will not be able convincingly to fill in the details. But as unskilled playwrights and actors know, even that will not suffice; to produce the desired end-effects, the system must also know the elements of drama, timing, and rhetoric, because just as orchestras need conductors and theatres need directors, all kinds of good performances require coherently organized styles. In other words, the best way to produce automated animation would be to program our machines in accordance with sound psychological theories. Today, few such theories yet exist, so it remains for us to invent them, to help the machine-based animators of the future direct their simulated screen-actors in ways that express and communicate.

But what is communication? This needs a theory, too — and I maintain that in order to communicate, one person must know (at least unconsciously) how to build structures and activate processes inside another person's brain. But because we cannot reach directly inside another person's brain, we have to work through the senses. How can automation help us use sensory experience to communicate ideas and goals, attitudes and dispositions, intentions and emotions? To dependably evoke the required types of subjective reactions, the animation engines of the future must be good psychologists!

But I do not mean to emphasize only the obscure and subtle aspects of this problem. For, first, we shall have to learn how to program our computers to accomplish the simplest-seeming, most commonsense goals — for example, that you must open the box before you put things in it, or that you must grasp and turn a doorknob before you can open a door. Present-day computers do not even know simple things like that! Then, once we've accomplished that, we must learn how to communicate to the audience more about our actors' goals. All animators understand that the procedures for making actors walk need not be based on precisely realistic descriptions of human locomotion gaits. More important is knowing how to shape those envelopes and trajectories, to express or indicate the picture-actor's mental state — its disposition in regard to determination, apprehension, expectation, disappointment, enthusiasm, and so forth. It is not enough that the picture-creature grasp a certain object; the film-maker must be able to express whether it was grasped by accident or intention, whether it was something the actor wants to possess or merely something to be moved out of the way while pursuing some more important goal.

How do we communicate such an idea?

We may be able to encode it into the trajectory itself, or in the actor's posture and apparent orientation; perhaps we can do it by focussing attention on the facial expression of the actor — or on the gestures of his antagonist. Every action has parameters to use for describing mental state. The more our machines can know about that, the more productive they can be.

I do not mean to say that we need to achieve complete scientific understanding of such matters. The great animators of the past accomplished miracles with little science and technology. But those talents were deplorably rare and the costs were too large for everyday use — and, instead, we could set as our future goal the development of technologies so powerful that the next generation of children will be able to use them fluently. Why, by the way, does animation so engage children? I suspect this is in part because expressive animation simplifies the audience's work — perhaps paradoxically, by immersing them in worlds that are so symbolic and abstract that the intentions and emotions of the objects being depicted are actually easier to apprehend and manipulate than the physical shapes and forms themselves!

We started with the common belief that scientists tend to use rational thought, while artists are intuitive. Yet when I discuss these things with my friends, I often find quite the opposite. Yes, scientists tend to be conscious and deliberate when they work on the details of their subjects, but they only rarely dwell upon the question of how they get their ideas — while our artists (if they are different at all) are less concerned with performances, but tend to reflect a good deal more about how to develop their concepts and skills.

(This essay is adapted from a talk given at Nikograf 89, a computer graphics conference held in Tokyo in November 1989.)