

This is the year 1300. Brother Giorgio, scholar-monk, has the task of making a map of Australia, a big island just south of India. Maps must record what is known about the places they represent, and Giorgio has been told about a strange Australian animal, ratlike, but much bigger, with a long thick tail and a pouch. He draws it, and it comes out like this:



Harold Cohen

Brother Giorgio's Kangaroo

Harold Cohen is a professor in the Visual Arts Department at the University of California at San Diego. In recent years he and AARON have been shown at the San Francisco Museum of Modern Art, the Stedelijk Museum in Amsterdam, the Tate Gallery in London, the Brooklyn Museum, the Ontario Science Center, and the Boston Science Museum. He has lectured widely on the subject of AARON, and AARON has reciprocated by providing Cohen with several thousand drawings, including designs for a series of mural projects. Cohen describes their relationship as symbiotic. Cohen is regarded as a pioneer in the application of AI to the visual arts.

A year later a world traveler is visiting Giorgio's monastery, and he tells our cartographer that he has the animal wrong. For one thing, it isn't *carrying* a pouch; the pouch is actually part of its belly. (Mercy! says Giorgio.) For another, it doesn't walk on all fours like a rat but on its hind legs, which are much bigger than its front legs. Giorgio redraws his picture:



But the tail rests on the ground. Giorgio tries once more. The traveler screws up his face in concentration, his eyes closed. I don't think that's quite right, he finally says, but I guess it's close enough.



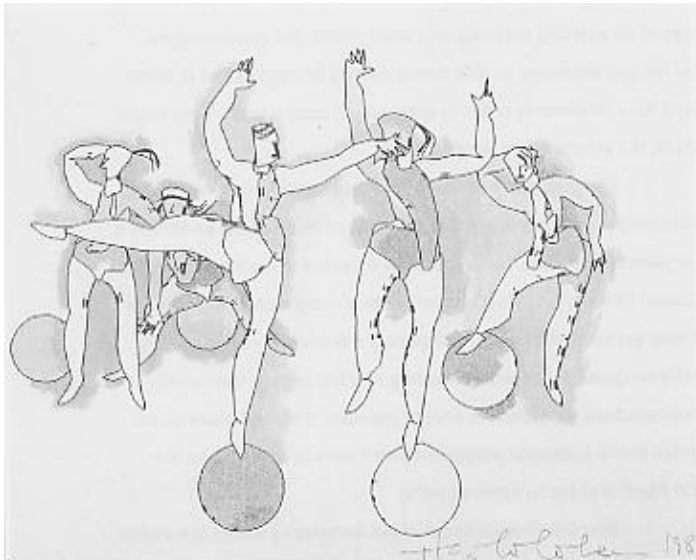
The year is 1987. AARON, a computer program, has the task of drawing some people in a botanical garden—not just making a copy of an existing drawing, you understand, but generating as many unique drawings on this theme as may be required of it. What does it have to know in order to accomplish such a task? How could AARON, the program, get written at all?

The problem will seem a lot less mystifying, though not necessarily less difficult, if we think of these two stories as having a lot in common. AARON has never seen a person or walked through a botanical garden. Giorgio has never seen a kangaroo. Since most of us today get most of our knowledge of the world indirectly and heavily wrapped in the understanding of other people from grade school teachers to television anchor persons, it should come as no surprise that a computer program doesn't have to experience the world itself in order to know about it.

How did Giorgio know about kangaroos before the visitor started to refine his knowledge? He had been told that the animal was ratlike, but how much good would that have done him if he had never seen a rat? For people, the acquisition of knowledge is cumulative, as it clearly has to be. Nothing is ever understood from scratch. Even the new-born babe has a good deal of knowledge "hard-wired" before it starts. And when we tell each other about the world, it isn't practical or even possible to give a full description of something without referring to some thing else. That's as true for computer programs as it is for people. There is an important difference, though. For people, knowledge must eventually refer back to experience, and people experience the world with their bodies, their brains, their reproductive systems, which computers don't have.

Harold Cohen, computer artist. (Photo by Lou Jones)





Athletes, a hand-colored, computer-generated drawing by Harold Cohen. (Photo by George Johnson)



From the *Bathers* series of hand-colored, computer-generated drawings by Harold Cohen. (Photo by George Johnson)

With this in mind, we might guess that AARON's knowledge of the world and the way AARON uses its knowledge are not likely to be exactly the same as the way we use what we have. Like us, its knowledge has been acquired cumulatively. Once it understands the concept of a leaf cluster, for example, it can make use of that knowledge whenever it needs it. But we can see what plants look like, and AARON can't. We don't need to understand the principles that govern plant growth in order to recognize and record the difference between a cactus and a willow tree in a drawing. AARON can only proceed by way of principles that we don't necessarily have. Plants exist for AARON in terms of their size, the thickness of limbs with respect to height, the rate at which limbs get thinner with respect to spreading, the degree of branching, the angular spread where branching occurs, and so on. Similar principles hold for the formation of leaves and leaf clusters. By manipulating these factors, AARON is able to generate a wide range of plant types and will never draw quite the same plant twice, even when it draws a number of plants recognizably of the same type. Interestingly enough, the way AARON accesses its knowledge of plant structure is itself quite treelike. It begins the generation of each new example with a general model and then branches from it. "Tree" is expanded into "big-tree/small-tree/shrub/grass/flower," "big tree" is expanded into "oak/willow/avocado/wideleaf" (the names are not intended literally), and so on, until each unique representation might be thought of as a single "leaf," the termination of a single path on a hugely proliferating "tree" of possibilities.

Obviously, AARON has to have similar structural knowledge about the human figure, only more of it. In part, this extra knowledge is demanded by AARON's audience, which knows about bodies from the inside and is more fussy about representations of the body than it is about representations of trees. In part, more knowledge is required to cope with the fact that bodies move around. But it isn't only a question of needing *more* knowledge; there are three different *kinds* of knowledge required—different, that is, in needing to be represented in the program in different ways.

First, AARON must obviously know what the body consists of, what the different parts are, and how big they are in relation to each other. Then it has to know how the parts of the body are articulated: what the type and range of movement is at each joint. Finally, because a coherently moving body is not merely a collection of independently moving parts, AARON has to know something about how body movements are coordinated: what the body has to do to keep its balance, for example. Conceptually, this isn't as difficult as it may seem, at least for standing positions with one or both feet on the ground. It's just a matter of keeping the center of gravity over the base and, where necessary, using the arms for fine tuning.

We started by asking what AARON would need to know to carry out its task. What I've outlined here constitutes an important part of that necessary knowledge, but not the whole of it. What else is necessary? Let's go back to Giorgio. Has it struck you that whatever Giorgio eventually knew about the relative sizes of the kangaroo's parts and its posture, he had been told nothing at all about its appearance? Yet his drawings somehow contrived to look sort of like the animal he thought he was representing, just as AARON's trees and people contrive to look like real trees and real people.

That may not seem very puzzling with respect to Giorgio. In fact, it may seem so unpuzzling that you wonder why I raise the issue. Obviously, Giorgio simply knew how to draw. I suspect that most people who don't draw think of drawing as a simple process of copying what's in front of them. Actually it's a much more complicated process of regenerating what we know about what's in front of us or even about what is not in front of us: Giorgio's kangaroo, for example. There's nothing simple about that regeneration process, though the fact that we can do it without having to think much about it may make it seem so. It is only in trying to teach a computer program the same skills that we begin to see how enormously complex a process is involved.

A hand-colored, computer-generated drawing of figures and trees with rocks in the foreground, by Harold Cohen. (Photo by Linda Winters)



How do humans learn to draw? To some degree, obviously, we learn about drawing by looking at other peoples' drawings. That's why we are able to identify styles in art, and why most of the drawings coming out of Giorgio's monastery would have had a great deal in common and be distinguishably different from, say, the drawings made in a Zen Buddhist temple in Japan. At the same time, all children make very much the same drawings at any one stage of cognitive development *without* learning from each other or from adults. They don't need to be told to use closed forms in their drawings to stand for solid objects, for example. That equivalent is universal; all cultures have used closed forms to stand for solid objects. In short, knowledge of drawing has two components. Giorgio learned about style, about what was culturally acceptable and what was not, from his peers. But before cultural considerations ever arise, drawing is closely coupled to seeing—so closely coupled that we might guess all major visual modes of representation in human history to have sprung directly from the nature of the cognitive system. So Giorgio never had to be told how to draw or how to read drawings. He could see.

He had to be told about kangaroos, not about how to draw kangaroos. Knowledge of drawing isn't object specific; if Giorgio

Black and White Drawing, a computer-generated drawing of figures and trees, by Harold Cohen. (Photo by Becky Cohen)



could draw a kangaroo, he could also draw an elephant or a castle or an angel of the Annunciation. If one can draw, then anything that can be described in structural terms can be represented in visual terms. That generality suggests that rather than thinking of knowledge of drawing as just one more chunk of knowledge, we should think of it as a sort of filter through which object-specific knowledge passes on its way from the mind to the drawing.

Like Giorgio, AARON had to be told about things of the world. Unlike Giorgio in having no hard-wired cognitive system to provide a built-in knowledge of drawing, it had to be taught how to draw as well, given enough of a cognitive structure (the filter just referred to) to guarantee the required generality. If provided with object-specific knowledge, AARON should be able to make drawings of those objects without being given any additional knowledge of drawing.

AARON's cognitive filter has three stages, of which the first two correspond roughly to the kinds of knowledge described above in relation to the human figure: knowledge of parts, articulation, and coordination. The third stage generates the appearance of the thing being drawn. Neither of the first two stages results in anything being drawn for the viewer, though they are drawn in AARON's imagination, so to speak, for its own use. First AARON constructs an articulated stick figure, the simplest representation that can embody what it knows about posture and movement. Then around the lines of this stick figure it builds a minimal framework of lines embodying in greater detail what it knows about the dimensions of the different parts. This framework doesn't represent the surface of the object. In the case of a figure, the lines actually correspond quite closely to musculature, although that is not their essential function. They are there to function as a sort of core around which the final stage will generate the visible results. Quite simply, AARON draws around the core figure it has "imagined." Well, no, not quite so simply. If you look at one of its drawings, it should be clear that the final embodying stage must be more complicated than I have said if only because AARON apparently draws hands and leaves with much greater attention than it affords to thighs and tree trunks.

AARON's embodying procedures are not like the preliminary edge-finding routines of computer vision, which respond to changes in light intensity without regard to what caused them. AARON is concerned with what it is drawing and continuously modifies the performance of this final stage with respect to how much knowledge has already been represented in the core figure. The greater the level of detail already present, the more AARON relies upon it and the closer to the core the embodying line is drawn. Also, greater detail implies more rapidly changing line directions in the

core, and AARON ensures a sufficiently responsive embodying line by sampling its relation to the core more frequently.

Nothing has been said here about how AARON's knowledge of the world is stored internally, about how its knowledge of drawing is actually implemented, or about its knowledge of composition, occlusion, and perspective. AARON's success as a program stands or falls on the quality of the art it makes, yet nothing much has been said about art and nothing at all about the acculturated knowledge of style, for which its programmer, like Giorgio's monastic peers, must admit or claim responsibility. All the same, there are interesting conclusions to be drawn from this abbreviated account. It should be evident, for example, that the knowledge that goes into the making of a visual representation, even a simple one, is quite diverse. I doubt that one could build a program capable of manipulating that knowledge and exhibiting the generality and flexibility of the human cognitive system other than by fashioning the program as an equivalent, artificial cognitive system. If nothing much has been said about art, it is because remarkably little of the program has anything to do with art: it constitutes a cognitive model of a reasonably general kind, and I even suspect that it could be adapted to other modes without too much distortion. But the lack of art specificity isn't as puzzling as it may seem at first glance. The principal difference between artists and nonartists is not a cognitive difference. It is simply that artists make art and nonartists don't.