HOW TO DRAW THREE PEOPLE IN A BOTANICAL GARDEN.

Harold Cohen

The University of California at San Diego

Department of Visual Arts, La Jolla, Ca 92093

Abstract

AARON is a program designed to investigate the cognitive principles underlying visual representation. Under continuous development for fifteen years, it is now able autonomously to make "freehand" drawings of people in garden-like settings. This has required a complex interplay between two bodies of knowledge: object-specific knowledge of how people are constructed and how they move, together with morphological knowledge of plant growth: and procedural knowledge of representational strategy. AARON's development through the events leading up to this recently-implemented knowledge-based form is discussed as an example of an "expert's system" as opposed to an "expert system." AARON demonstrates that, given appropriate interaction between domain knowledge and knowledge of representational strategy, relatively rich representations may result from sparse information.

Figure 1: AARON drawing, 1987

1 Preamble

Brother Giorgio is a 12th Century scholar-monk whose task it is to record what is known of the world's geography, and he is currently making a map of Australia, an island just off the coast of India. Since an essential part of map-making involves representing the animals of the country, he is making a drawing of a kangaroo. Now Brother Giorgio has never seen a kangaroo. But he understands from what he has been told that the kangaroo is a large rat-like creature with a pouch, and with an exceptionally thick tail. And he draws it accordingly (figure 2a).


Figure 2a

While he is so engaged, a traveller visits the monastery, and he tells Giorgio that his drawing is wrong. For one thing -- and Giorgio finds this quite implausible -- the kangaroo doesn't carry a pouch -- its pouch is part of its belly! And, says the traveller, it doesn't go on all fours: it stands upright, on rear legs much bigger and thicker than the front legs (figure 2b).


Figure 2b

And the tail doesn't stick straight out, it rests on the ground. Giorgio completes all the necessary changes, and the traveller assures him that though he hasn't got it quite right, it's close (figure 2c).

AARON, late in the 20th Century, is a knowledge-based program that is capable of the autonomous generation of original "freehand" drawings, like the one in Figure 1. Like Brother Giorgio, AARON has never seen the things it draws.


Figure 2c

It, too, is obliged to rely upon what it is told. Unlike Giorgio, however, it cannot make use of associative knowledge. There would be no point in telling it that a kangaroo looks a bit like a rat, for example, not only because it doesn't know about rats, but because it has never looked at anything. What both Giorgio and AARON make clear is that the plausibility of a representation does not rest upon the correctness of the knowledge it embodies. Indeed, for anyone lacking knowledge of marsupials, the "correct" knowledge of the kangaroo's pouch is at least as implausible as Giorgio's initial understanding. Nor does plausibility rest upon the completeness of that knowledge, since representations only ever represents what is being represented with respect to an arbitrarily small set of properties. Given one important proviso -- that the representation-builder has general knowledge about how to make representations -- there would appear to be no lower limit of knowledge below which the making of a representation is impossible.

And that proviso points to the main thrust of this paper. It will show AARON's visual representations to involve a spectrum of representational procedures, and a spectrum of different kinds of world knowledge. It will also show the degree to which the particular quality of those representations depends upon the intimate meshing of the program's world knowledge with its knowledge of representing.

AARON has been under continuous development for nearly fifteen years now and it has gone through many generations. At fifteen it may well be the oldest continuously-operational expert system in existence and perhaps the only computer program to have had something akin to a life-story. But perhaps AARON would be better described as an expert's system than as an expert system: not simply because I have served as both knowledge engineer and as resident expert, but because the program serves as a research tool for the expansion of my own expert knowledge rather than to encapsulate that knowledge for the use of others.

The goal of this research is to understand the nature of visual representation. The term should not be understood to imply the various mechanical methods -- perspective, photog-

raphy, ray-tracing -- by which two-dimensional transforms of a three-dimensional world may be generated. All of these are knowledge-free, in the sense that the photographic process neither knows, nor needs to know, what is in front of the lens. AARON helps to address questions that are both more fundamental and more general. What do computer programs -- and, paradigmatically, human beings -- need to know about the external world in order to build plausible visual representations of it. What kind of cognitive activity is involved in the making and reading of those representations?

The making of representational objects -- the drawings, paintings, diagrams, sketches in which representations are embodied -- constitutes the only directly-examinable evidence we have of "visual imagining." I mean those internal cognitive processes that underpin and inform the making of representational objects, and which we all enjoy to some extent, whether or not we make representational objects. I assume that the reading of representations involves essentially similar processes. But making requires more than reading does. It requires a special body of knowledge -- knowledge of representation itself -- that is part of the expertise of the artist, just as the representation of a body of knowledge within an expert system requires an analogous expertise of the knowledge engineer. Understanding the nature of visual representation requires asking what artists need to know in order to make representational objects; what they need to know, not only about the world, but also about the nature and the strategies of representation itself.

AARON's role in this investigation, then, has been to serve as a functional model for a developing theory of visual representation. The stress is on the word "functional," for the most convincing test of a theory of representation is the model's ability to make representational objects, just as the plausibility of a theory of art resides in art-making.

AARON was last reported in detail in 1979, in the proceedings of IJCAI-6, at which time it was making drawings like that in figure 3. The differences in its output have been matched, of course, by large changes in the program itself. But these have been developmental rather than radical changes, following a pattern analogous to that of human cognitive development, and AARON has retained its identity and its name throughout.

Figure 3: AARON drawing, 1979

Part of my purpose here is to describe the current state of the program. The other part is to account for its development. That means, necessarily, to describe the domain of interaction between program and programmer, to delineate the purpose that the one serves for the other. AARON has been a research tool for me, but also something very like an artist's assistant, capable always of enacting, without human aid or interference, the understanding of art embodied in its structure. And my relationship to the program has become increasingly symbiotic. Without AARON's sophisticated enactment of my own understanding, that understanding would not have developed as it did. Without that developing understanding AARON could never have become the sophisticated adjunct artist that it is.

My long-held conviction that AARON could only have been written by a single individual has been based on rather vague suspicions of cultural incompatibilities existing between the disciplines of knowledge engineering and art. Now I believe, rather more precisely, that the problem -- and, indeed, a fundamental limitation of expert systems -- lies in the artificial separation of two bodies of knowledge, that of domain-expert on the one hand and knowledge-system-expert on the other.

2 Aaron: Early Versions

In all its versions prior to 1980, AARON dealt with exclusively internal aspects of human cognition. It was intended to identify the functional primitives and differentiations used in the building of mental images and, consequently, in the making of drawings and paintings. The program was able to differentiate, for example, between figure and ground, and insideness and outsideness, and to function in terms of similarity, division and repetition. Without any object-specific knowledge of the external world, AARON constituted a severely limited model of human cognition, yet the few primitives it embodied proved to be remarkably powerful in generating highly evocative images: images, that is, that suggested, without describing, an external world [Cohen, 1979]. This result implied that experiential knowledge, inevitably less than constant across a culture and far less so between cultures, is less a determinant to the communicability of visual representations than is the fact that we all share a single cognitive architecture.

From the program's inception around 1973, I had been convinced that AARON would need to be built upon a convincing simulation of freehand drawing, and gave much attention to modeling the feedback-dependent nature of human drawing behavior. As a consequence of this stress the program was formulated, initially, largely in terms of line generation. Closed forms, those universal signifiers for solid objects, also were generated from rules directing the development of lines: rather like the way one might drive a closed path in a parking lot by imagining a series of intermediate destinations, veering towards each in turn and finally returning to one's starting point [Cohen, Cohen, Nii, 1984]. Following a paradigm we see exemplified in rock drawings and paintings all over the world, AARON observed a general injunction against allowing closed forms to overlap each other, and would be obliged to modify its closure plans frequently in order to prevent over-

lap. This resulted in a richer, less predictable, set of forms than the unmodified closure rule would have permitted. But underlying this richness was the fact that AARON had no prior spatial conception of the closed forms it drew: their spatial presence, their identity, was the result, not the cause, of an entirely linear operation.

Throughout this phase of AARON's development, a constant sub-goal was to increase the range and diversity of its output. And in 1980 this desire led to the development of a new generating strategy for closed forms. It had its basis in an attempt to simulate the drawing behavior of young children, specifically at that immediately post-scribbling stage at which a round-and-round scribble migrates out from the scribble-mass to become an enclosing form (figure 4). It was while this work was in progress that a colleague expressed an interest in having AARON make "realistic," as opposed to evocative, drawings. Could it, for example, make a drawing of an animal?


Figure 5a: AARON, animal drawing

Figure 4

Figure 5b: AARON, animal drawing

I must avoid here what would be a lengthy digression on the nature of realism. Let it suffice to say that I took my colleague's words to imply a visual representation of an animal, as opposed, say, to a diagram. Since I had never drawn animals and had little idea about their appearances I thought it unlikely that I could oblige. What little knowledge I could place at AARON's disposal was barely sufficient to construct a diagrammatic stick figure: a representation, certainly, but not a visual representation.

Now it happens that the "enclosing" stage of children's drawing is also the stage at which they begin to assign representational significance to their drawings. If this was more than a coincidence, I speculated, perhaps it would be possible to generate an adequate representation by enclosing a stick figure the way a child encloses a scribble. It proved to be a good guess. On the first attempt the program's drawings showed a startling resemblance to the rock drawings of the African Bushmen (figure 5a). Encouraged by the result I amplified the program's knowledge to take some account of the bulk of the animal's body, and the drawings shifted their stylistic affiliations to the caves of Northern Europe (figure 5b) [Cohen, 1981].

In retrospect it seems obvious that the closed forms of these drawings would have produced a richer evocation of "real" animals than a diagrammatic stick-figure could. The clear differentiation of style that resulted exclusively from the change in the enclosed, subsequently invisible, diagram is more problematic, however. Perhaps "style" in art is less a question of autography than of what the artist believes to be significant.

AARON was now potentially able to generate a large variety of geometrically complex closed forms without requiring geometrically complex descriptions. The gain was obvious enough to ensure to this new strategy a permanent place in AARON's repertoire, even without the goal of visual representation. From that point forward, all closed forms involved the initial construction of a "conceptual core," corresponding to the child's scribble, around which the form was "embodied" (figure 6).

One important result of the new strategy was to shift the stress in AARON's drawing mode away from its initial linearity, yet the greater gain had less to do with the growth of AARON's formal skills than with its "cognitive" development. For the first time AARON now had some concept of what it was to draw before it began to draw it.

Which did not mean that AARON proceeded to draw real-world objects; on the contrary, the representation of real-world objects seemed as unnecessary to my research goals as it was inconsistent with my own history as an artist. The animals disappeared from AARON's repertoire and no further attempt was made at that time to apply the new strategy to the representation of real-world objects. Yet even in the absence of real-world knowledge, the new cognitive mode endowed AARON's images with an increasingly "thing-like" presence that seemed to demand an explicitly visual space in which to exist. Thus, for example, where the earlier versions of the program had avoided overlapping figures, occlusion now became a fundamental principle of pictorial organization. By 1984 the earlier "rock-art" pictorial paradigm had given way entirely. The pressure to provide real-world knowledge of what AARON's new visual space contained became inescapable and the first of several knowledge-based versions of the program was constructed (figure 7).

Figure 6: AARON drawing, 1983

Figure 7: AARON drawing, 1985

I do not intend by this account to imply some metaphysical force guiding AARON's development and my own hand. Nor is it necessary to do so. Every system has its own logic, and the need to follow the dictates of that logic, to discover where it will lead, may easily transcend the private inclinations of the investigator.

3 AARON: Recent and Current Versions

I said earlier that the goal of this research is to discover what the artist needs to know about the world in order to make plausible representations of it: not correct representations, or complete representations, or realistic representations -- none of these notions hold up very well under examination -- but plausible representations. If I had asked *how much* the artist needs to know, the answer would have been that the question is hardly relevant: we make representations with whatever we know. Given adequate knowledge of the representational procedures themselves, there is virtually no lower limit of world-knowledge below which representation is impossible. The goal, rather, is to discover *how* representational structures represent *what* they represent: how we use what we know to build those structures.

What does AARON represent, and how -- by means of what structures -- is it represented?

As the title of this paper suggests, AARON represents a small part of the flora and the fauna of the world, with a little geology thrown in: a tiny part of the whole of nature. Because plausibility does not rest upon how much the image-maker knows about the world AARON has never been provided with a large object-specific knowledge base -- large, that is, in the sense of referring to many different objects. And because object-specific knowledge is also purpose-specific, no attempt has been made to give it knowledge that might be considered essential for representations of other kinds than its own and within other disciplines. Most particularly, its object-specific knowledge contains very little about appearances, and the program's overall strategy rests upon being able to accumulate several different kinds of non-visual knowledge into visually representable forms. This is not a neatly sequential process. As I will show, different knowledge is called into play at different stages of accumulation; the program's representational knowledge is not simply invoked as a final step.

In the category of object-specific knowledge the program has five levels, each with its own representational formalism. At the first level is AARON's declarative knowledge. For example: a human figure has a head, a torso, two arms and two legs. A plant has a trunk, branches and leaf clusters. This declarative knowledge is represented outside the program itself in frame-like forms that are read in as they are needed. So, also, is knowledge of several pictorial "classes." A class is characterized simply by what elements may be used in a given drawing and -- since AARON does not use an eraser -- the front-to-back order in which they may be used. Thus AARON begins a new drawing by selecting a pictorial class, and proceeds by expanding each entry in the class heirarchically into an internal tree-structure, at the lowest levels of which are the management procedures responsible for the generation of individual elements of the drawing. There is, for example, a

"hand" manager whose sole task is to produce examples of hands on demand, to satisfy the specifications that are developed for it.

The expansion of externally-held declarative knowledge into internal tree structure is done on a depth-first basis, and AARON does not know in advance what the current class will require at a later stage; and it may, in fact, over-ride the demands made by the class in favor of constraints that have developed within the drawing. A class is only minimally prescriptive; it will call for "some" trees or people, rather than two trees or three people, where "some" may be specified, for example, as any number between zero and four. Consequently the expansion is not deterministic. Decision-making is relatively unconstrained at the start of the drawing and, though it becomes increasingly constrained as the drawing proceeds, AARON randomizes unless it has some clear reason for preferring one thing or one action over another, as people do. All higher-level decisions are made in terms of the state of the drawing, so that the use and availability of space in particular are highly sensitive to the history of the program's decisions.

AARON's first and ongoing task, then, has to do with the disposition of its few objects in a plausible visual space.

3.1 The Nature of Appearances

When I first provided AARON with the knowledge it would need to make blatantly representational drawings, I reasoned that, since anything one sees through a window is as real as anything else, pictorial composition was hardly relevant to the issue of plausibility. I assumed, therefore, that I could safely fall back upon the simplest, and perhaps the most universal, of compositional paradigms: put it where you can find space for it. And this paradigm, extensively used in AARON's two-dimensional days, remained valid in its new world to the extent that three people in open view make neither a better nor a worse composition than five people hiding in the foliage. A fundamental problem emerged, however, centered on the ambiguity of the word "where." Until recently AARON has never had a fully 3-dimensional knowledge-base of the things it draws: foreshortening of arms or the slope of a foot in the representation were inferred from AARON's knowledge of the principles of appearance, not by constructing the figure in 3-space and generating a perspective projection. And it happened too frequently in the program's first efforts at representation that people in the picture would stand on each other's feet (figure 8).

I've been using the term "plausible representations" to mean representations that are plausible with respect to appearance, and I must now consider what appearance means and what it implies. Appearance implies what the world looks like. It implies the existence of a viewer, and a viewpoint that controls the disposition of objects within the viewer's visual field. Since much of what the viewer sees is illuminated surfaces, it implies also some condition of lighting that controls visibility in some particular way. And since lighting is arbitrary with respect to the object itself it follows that the appearance of objects -- as opposed, for example, to their structure, their mass or their dimensions -- is a transitory characteristic. In order for appearance to imply specific knowledge of how particular objects look under particular and transitory lighting


Figure 8: AARON drawing, 1986

conditions, we would have to be able to store and retrieve, not merely "visual fragments," but complete "mental photographs." And that is surely not the general case.

On the other hand, we can regard the way solid objects occlude each other, the way objects take less space in the visual field as they get further away, the way light falls on surfaces and so on, as a set of principles. In theory we should be able to infer a particular appearance by applying the principles of appearance to a particular surface description; that is exactly what the various strategies of "solids modeling" do. But the human mind is rather poor at inferring appearance, partly because it rarely has adequate surface descriptions available to it -- we use appearance to provide those descriptions, not the other way around -- and partly because the human cognitive system makes use of a gamut of "cognitive perspectives" quite unlike the self-consistent geometries upon which solids modeling relies. One result is that in the one period of history when art has concerned itself explicitly with appearance -- the western world since the Italian Renaissance -- it has inferred the appearance of simple surface configurations, but has relied heavily upon direct observation for the depiction of complex surfaces. For example, the artists of the Renaissance used perspective in depicting objects with simple surfaces -- buildings, tesselated floors -- but almost never attempted to use perspective in depicting the human figure (figure 9). And, of course, solids modeling has balked at the surface complexity of the human figure for the same reason: the difficulty of providing adequate surface descriptions.

Figure 9: Fra Angelico, Annunciation, 1437

3.2 Pictorial Organization versus Dimensional Plausibility

Fortunately, the cognitive system provides a convenient short-hand for the representation of surfaces. Since the eye functions as a contrast amplifier we are able to detect the bounding edges of surfaces very efficiently, and we make heavy use of the behavior of those edges to provide information about the surfaces inside them. In using edges as the basis for a representational mode, then, much of the problem of surface illumination is bypassed. Plausibility rests upon the behavior of the edge, and upon issues that can be addressed in terms of edges: notably occlusion and spatial distribution.

Actually, very little is required, in terms of occlusion and perspective, in drawing a single figure or a single plant. However, the need to represent objects plausibly with respect to other objects requires a significant level of control over 2-space placement and the relative sizes of objects within the representation, and requires more extended control of occlusion.

This is more complicated than it may seem. As I remarked, visual representation in general rests upon a complex of cognitive "perspectives," not upon the automatic and knowledge-free 2-space mapping of the visual field provided by photography, or its computer-based equivalent, solids modeling. So great is the difference between the cognitive view and the automated view that experienced studio photographers habitually use polaroid instant film while setting up a shot in order to find out what the world they can see clearly in three dimensions will look like in two. Evidently 2-space organization cannot be adequately predicted or controlled exclusively through control of the 3-space arrangement of objects. Nor, conversely, is it possible to guarantee plausibility with respect to physical dimensionality by concentrating exclusively on pictorial organization.

In constructing visual representations the human artist appears to work under two simultaneously-active and mutually-exclusive constraint-sets. The "imaginational planning" that marks this mode is best evidenced by the artist's sketch, in which 2-D space is allocated less to the objects to be represented than to the space they occupy, and in which space is increasingly committed as those objects "congeal" simultaneously within both their actual 2-D space and their referenced, implied 3-D world (figure 10).

Figure 10. Eugene Delacroix 1862 "Resurrection".

At present, AARON uses only a crude, static model of this essentially dynamic process. While it organizes primarily in 2-space terms, it also maintains a floor plan of the "real" world it is depicting. Space for an object is allocated initially on the plane of the representation. It is then projected back into the "real" world, where it is adjusted to ensure valid 3-space placement, and then it is projected forward again into the representation. In doing this, perspective is used only to the degree of determining where the bases of objects -- the feet of figures -- will fall, and how high the objects will be, in the representation. It thus ensures that real-world objects are placed plausibly with respect to their ground-plane while doing very little about planning in 3-space terms. People no longer stand on each others' feet, but a genuinely dynamic model of this imaginational planning remains a goal for the program's future development.

3.3 Levels of Knowledge

During the expansion process, a second level of knowledge -- exemplary knowledge -- is invoked to provide fuller specification for the management procedures. The determination that this particular figure will have a large head and long arms, for example, involves applying the descriptors "large" and "long" plausibly to a set of prototypical dimensions held in table form within the program. The further determination that this figure will hold a particular posture, requiring its right arm to be extended horizontally and its right hand to be pointing, will require three further levels of amplification before an adequate specification can be generated. First: the figure is articulated, and AARON has to know where the articulations are (structural knowledge). Second: it must know what the legal range of movement is at each articulation (functional knowledge). Third, and most important, since a coherently-

articulated figure is more than a random collection of legal movements, there has to be knowledge of how a figure behaves; how it keeps its balance and how it gestures.

AARON's knowledge of plants follows similar patterns of distribution. AARON understands plant morphology in terms of branching, limb thickness with respect to length and branching level, the clustering patterns in leaf formations, the size of the plant, and so on. It has no stored descriptions of particular plants, and its entire plausible flora is generated by the same small set of management procedures, through the manipulation of these morphological variables (figure 11).

3.4 From Stick Figure to Solid Figure

Note, however, that these specifications contain no reference to appearances, and that they suffice only to inform the production of plausible stick-figure representations. Where the expression of object-specific knowledge spans a range of forms from conceptual to dimensional, the expression of visually representational knowledge requires visual, two-dimensional terms, and the stick-figure, is, in fact, the first visual manifestation of AARON's amplified knowledge. All postural issues are determined in relation to it alone.

Figure 11: AARON drawing, 1987

The organization of AARON's object-specific knowledge is thus a five-tiered structure in which each successive level is accessed at the appropriate time for what it can add to the whole. Broadly speaking these levels span a spectrum of knowledge types from wholly declarative and external to wholly procedural and internal, and the program proceeds from the general, where it manipulates conceptual tokens like "hand," "pointing," "large," "some," to the specific, detailed and plausible instantiation of these tokens.

In the second stage, the stick-figure is used as an armature upon which to build a more extensive framework. This is the stage at which the exemplary knowledge of the thickness of the parts is invoked. The lines of this framework do not represent the external surfaces of the figure. They are loosely associated with musculature and skeletal features -- for example, the single line representing the hip-to-hip axis in the stick figure is expanded into a diagrammatic pelvis -- but their primary function is to guarantee sufficient bulk to the figure in

whatever posture, and from whatever position, it is viewed. With the completion of this stage AARON has provided itself with the conceptual core of its representation, similar functionally to the young child's scribble. And it is around this conceptual core, in the third and final stage, that the figure is embodied (Figure 12 shows an incomplete core taken from a current, fully 3-D version of the program).


Figure 12: partial core figure

Embodying involves generating a path around each of the parts of the conceptual core. These are taken, as the elements of the drawing are, in closest-first order. Part of the internal representation of the drawing that AARON maintains for itself consists of a matrix of cells onto which are mapped the lines and the enclosed spaces of the drawing. Thus the conceptual core is now recorded as a mass of marked cells, to develop a path around which AARON uses what is, in essence, a simple maze-running procedure. However, its implementation rests heavily upon the fact that AARON draws, as the human artist does, in feedback mode. No line is ever fully planned in advance: it is generated through the process of matching its current state to a desired end state. As with any feedback-controlled system, AARON's performance is characterized by its sampling rate and by how radically it corrects. This part of the program most intimately determines AARON's "hand", and it has not changed greatly since the program's earliest versions.

Unlike the earlier versions, however, the strategy for "imagining" the intermediate destinations around its path depends upon two things. Firstly, upon its ability to recognize and to deal with a number of special-case configurations in the core figure (figure 13a, b). These -- and most particularly a configuration indicating a sharp concavity -- are intimately

involved in the self-overlapping folding of outlines that convey so much about the appearance of complex three-dimensional forms. Secondly, AARON knows what it is drawing, and it associates some particular degree of carefulness with the delineation of any particular element. This knowledge is expressed in the use of an additional feedback parameter: the distance from the core at which the path will be developed. Thus, for example, it will draw a thigh rather loosely -- that is, at some distance from the conceptual core and with a relatively low sampling rate -- while it will draw a hand close to the core and with a high sampling rate. Both of these are controlled by the placement and the frequency of the intermediate destinations around the marked-cell mass. AARON further adjusts its own sampling rate and correction with respect to the size of the element it is drawing relative to the size of the entire image.


Figure 13a: strategy for concave configuration

Figure 13b.   Edouard Manet 1862 "Study for a Woman at her Toilet".

4 Conclusion

In practice AARON makes drawings of whatever it knows about without requiring any further instructions for the making of a particular drawing -- and, indeed, without possessing any mechanism through which it could take instructions. To the degree that it does nothing much more than enact what it knows, AARON provides an intuitively satisfying model of "visual imagining," in that it permits the expansion of relatively sparse real-world, object-specific knowledge into a convincing representation of a visual experience.

I have described AARON's knowledge as falling into two broad categories: what it knows about a small range of world objects and what it knows about building visual representations. And I have proposed that these two categories must be intimately inter-related in any satisfactory model of human knowledge-based performance. The conclusion is an obvious one; we can only represent what is representable in terms of available representational strategies. I have no doubt, for example, that the program's development has been profoundly determined by the fact that it has been written in 'C' rather than in LISP. AARON's representational strategy, deriving as it does from the young child's relatively undifferentiated perceptions of the world, is well adapted to the representation of blob-like forms, or forms with a strong axis -- heads and limbs, for example. Yet AARON is unable to deal with cube-like objects, the perception of which rests upon high contrast edges in the center of a form as well as at its extremities. AARON will need new representational strategies, not merely more object-specific knowledge, before it can present a new view of the world, just as the young child is obliged to develop new strategies before it is able to put the sides on its representations of houses (figure 14).


Figure 14: childrens' drawings

Finally: I have claimed for AARON only that it makes plausible representations, and have left aside the consensus judgement that its drawings represent a high level of artistic accomplishment. Why have I had nothing to say about "aesthetic" principles like harmony and balance?

The short answer is that AARON is entirely unaware of the existence of those principles, and that since its drawings are aesthetically satisfactory, we must surely question the relevance of those principles to artistic production. This is not to say that AARON does not embody principles of its own, but that whether these are aesthetic principles is largely a matter of definition. I have to assume that the simple "find-enough-space" rule to which I referred earlier contributes to the aesthetic appeal of the outcome, but it is quite different in kind from the aesthetic rules commonly believed to guide the production of works of art.

The fuller answer is that I regard "style" -- surely the most difficult word in the entire vocabulary of art -- as the signature of a complex system. I regard the aesthetics of AARON's performance as an emergent property arising from the interaction of so many interdependent processes, the result of so many decisions in the design of the program, that it becomes meaningless to ask how much any one of them is responsible for the outcome. If AARON has maintained a consistent aesthetic, a consistent identity, from its earliest endeavors, I have to assume it to reflect consistent patterns of my own in determining its development. If someone else wrote a similar program I would expect it to exhibit a different identity and a different aesthetic.

That answer would be begging the question, if the point of the question was to consider how an orthodox expert system might be built to generate objects of high artistic value. That isn't the point; given the orthodox separation of domain knowledge from representation knowledge, I do not believe it will be in the foreseeable future. This is one place where it seems not to be true that two heads are better than one.

References

[Cohen, 1979] Harold Cohen. *What is an Image?* Proceedings of IJCAI-6, Tokyo, 1979.

[Cohen, 1981] Harold Cohen. *On the Modeling of Creative Behavior.* Internal Paper for the Rand Corporation. 1981.

[Cohen, Cohen, Nii, 1984] Harold Cohen, Becky Cohen and Penny Nii. *The First Artificial Intelligence Coloring Book.* William Kaufman, 1984.

Technical Note

While the earliest versions of AARON were built as production systems, all the more recent versions have been strongly object-oriented, as the above discussion might indicate. The program has about 14,000 lines of 'C' code and occupies almost a half-megabyte of compiled code, exclusive of external declarative knowledge structures and the internal representations of the developing drawing it makes for its own use. The most recent version was written under UNIX on a MicroVax-II, on which machine a single drawing takes about an hour of CPU-time, and has been ported to several other UNIX machines.