

THE MATERIAL OF SYMBOLS

First. Annual Symposium on
Symbols and Symbol Processes

UNIVERSITY OF NEVADA
LAS VEGAS

Harold Cohen
Center for Art/Science Studies
UNIVERSITY OF CALIFORNIA AT SAN DIEGO
August 1976

Karin and Sherry are seven year-old twin sisters. They are both in the habit of carrying large bags of colored pens and pencils with them wherever they go, and at every possible opportunity they sit down on the floor and start to make drawings. Those illustrated her (1,2) are entirely typical of their work. Their output is certainly well above average in quantity, but the drawings themselves are in no major respect atypical of the sort of drawings which most Western children might make at some period of their development.

On the same afternoon that these drawings were made, I proposed a game to them: I would cover sheets of paper with dots, and they would make their drawings by joining up the dots (3»~)- They both took to the game with obvious enjoyment, but also with an unexpected attention to the structural constraints imposed upon their performance by these new rules, which they promptly investigated. One of them wanted to know whether it was permitted to leave some of the dots unused. The other asked whether she was allowed to use the dots as eyes, if she was drawing a face, and in due course she contrived to use the dots also as Christmas tree decorations, snowflakes, sunbeams, and a number of other unspecified objects which she said were falling from the sky (3). In both cases - and I do not believe that this is part of their normal procedure - each drawing was followed by a long verbal account of the subject matter. After an hour or so the game ended, and they returned to their habitual mode.

There are a number of formal differences in the drawings which result from the two modes which might be dealt with at some length. Their normal practice, for example, is to use the plane of the paper to represent some sort of coherent spatial unity, corresponding very roughly to what we might call a "view". In some of the dotted drawings this practice gives way to a more elemental approach, in which the plane of the paper is used in a manner largely neutral with respect to the images, and the images are disposed upon it without regard to any concept of "natural" ordering in the real world. (By the way, these results are quite consistent with the results of a more extensive set of similar experiments with a drawing class at UC San Diego.

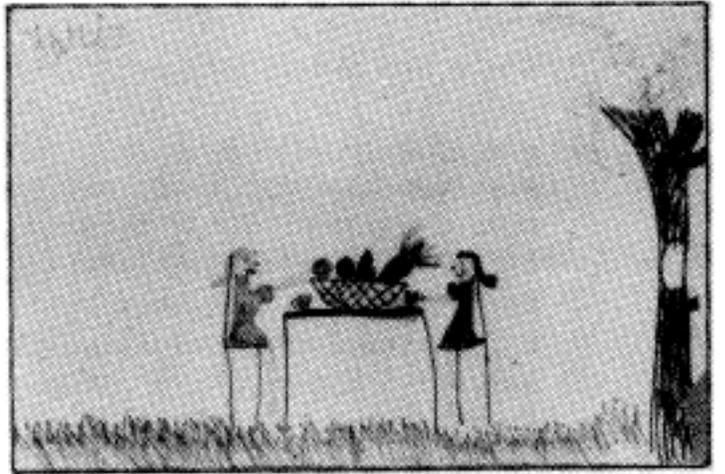
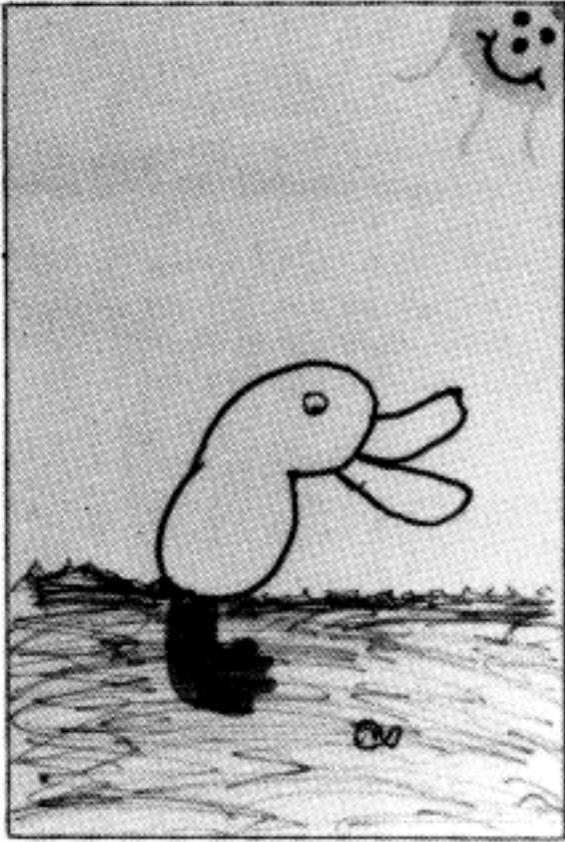
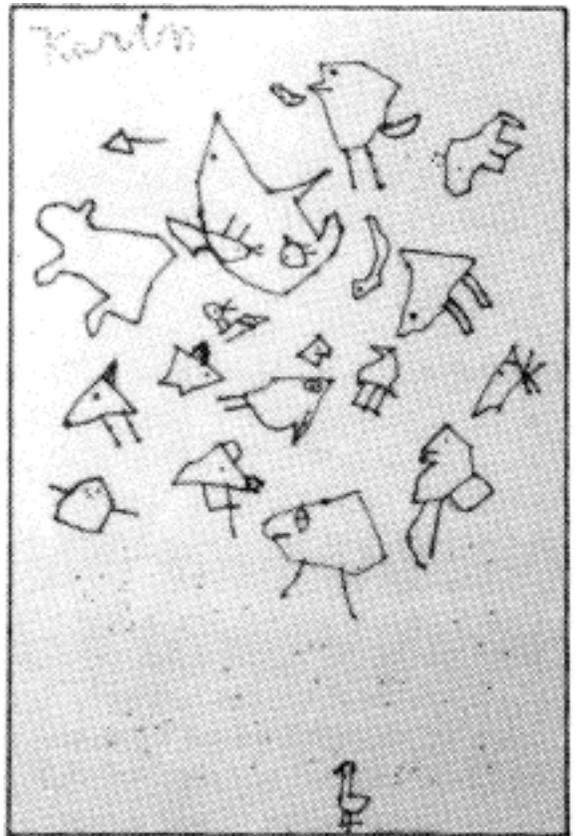
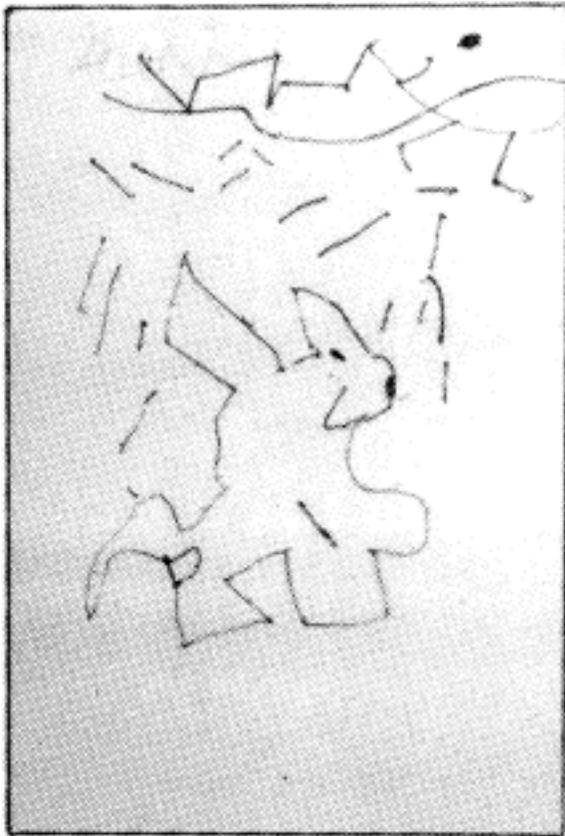


fig 2

fig 1

fig 3

fig 4



The students there were up to forty years older than Karin and Sherry, and their habitual modes involved different conventions to those of the children; but they were certainly no less conventional.)

But the more immediately noticeable differences between the two modes relate less to the formal aspects of the drawings than to the level of imagination and inventiveness which Sherry and Karin exercise in making them. When Karin decides to sign her name in a manner appropriate to the game (4), she is making a witty comment about the nature of drawing at a level of insightfulness we might not expect from a seven year-old. If we compare the bird in one of her dot-drawings (5) with the drawing of a duck made just half an hour earlier, we are struck by the fact that she is evidently capable of rather acute observation, although it required the setting up of unfamiliar, and presumably challenging, circumstances to allow her to exercise that capability.

What becomes clear, in fact, is that there is a significant difference between an image of a bird, and an image of an image of a bird. The earlier drawing is less a duck than it is a toy duck, less the result of observing what the real world is like than it is the result of learning what drawings - of the world - are supposed to look like. It is conventional in the precise sense that its conventions are the common property of that sub-culture we call children, where their stability is maintained both by the children's desire to conform and by the adult desire that they should.



fig 5

I incline strongly to the view that we all spend our lives -- not merely our childhood -- trying to effect an acceptable and workable compromise between the internal demands for the satisfaction of our individual psychic needs, and the demands made upon us by the culture within which we live, for the sake of the stability, if not necessarily the ultimate well-being, of the culture itself. This is not to say that the things we do, like drawings, singing, talking,

do not grow from the most fundamental patterns of the mind but that the cultural rules imposed for their exercise may lead to behavioral patterns quite at variance with these deeper ones.

Most children are able to build their early images without difficulty with marks which result directly from simple physical movements, just as the African sculptor has no difficulty satisfying his representational needs with conceptually simple forms requiring simple manufacturing skills (6).

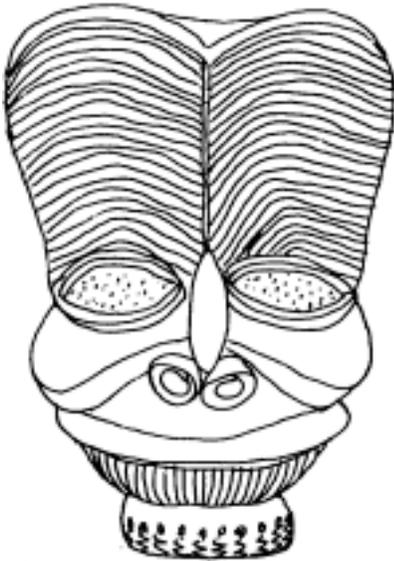


fig 6

The notion of representation which held sway in Europe for nearly five hundred years, on the other hand, requires the student to spend a minimum of three years persuading his eye to see what it is supposed to see, and disciplining his hand to move as it is supposed to move. These movements are arbitrary with respect to the individual, since they have to be determined by events in the world -- the random play of light and shade on objects -- which have nothing to do with the way his shoulder and wrist are articulated. The reconciliation which the artist in this tradition is obliged to make is a striking example of the sort of compromise I am referring to.

We do not pay for our membership of the culture on a one-day-on, one-day-off basis. All our behavior is acculturated to some degree, and any attempt to isolate a discreet behavioral mode which we might think of as "natural" would be fruitless. Yet we might still find in the underlying structures of behavior aspects which are evidently not fashioned by the constraints of any particular culture, and this would be as close as we might come to a notion of "naturalness". It will be the tracking down of these aspects with which I will be concerned, knowing very well that their separation from other aspects is a theoretical one.

Much of our mental activity seems to involve complex schema

of entities standing for other entities, and we would probably agree that the externalizing and manipulation of images, as such, grows directly from basic mind functions. But that area of symbol-manipulation which is directed towards communication between individuals and between groups must obviously involve highly acculturated performance. For a symbolic structure to stand any reasonable chance of being unambiguously understood, its maker must both have clear knowledge of the expectations which the reader will bring to its reading, and be prepared to accept the constraints imposed by those expectations. Communication is possible within a culture only because of existing agreements as to what entity is to stand for what entity, and how it is to be presented to be recognized as doing so. At an even more basic level, this implies also that all the involved parties know about the same entities: which may be true, more or less, within the same culture, but is unlikely to be true from one culture to another.

These would seem not to be very promising conditions for the exercise of imagination, inventiveness, and all those other virtues we associate with the making of art, or, indeed, for our understanding of art produced by any culture other than our own. But I think we have to conclude that art never has been devoted primarily to the cultural function of communication, and indeed it may never have been thought that it did before our own time. The more historic view within our own culture pictures the artist in communion with variously-conceived extra-human sources of inspiration and wisdom,, explicitly acknowledging the fact that if he speaks on behalf of the community, he does not speak with its voice or in terms which will necessarily be understood.

Art history deals with the problem of tracking and identifying the transformations which continuously modify the significance of symbols within the changing cultural continuum. But there are other problems of a more fundamental kind which fall outside the scope of orthodox iconology. Any art theory which begins with a view of the artist as serving primarily the cultural need to formulate and transmit explicit meanings inevitably ends up viewing the whole system as a sort of noisy telephone network, in which

the receiver strives constantly to reconstruct the original message. Yet the cultural mismatch between artist and viewer must then be a major source of noise in the system, and we account for the discrepancies between what the artist "has in mind" and what the viewer thinks he understands, by the notion of "interpretation". We do not necessarily have any evidence beyond our own "interpretations", however, as to what, if anything, the artist had in mind in the first place.

This emphasis upon the specifically cultural use of symbols has left us without any account of the underlying structures of image-generating behavior more convincing than the Divine Muse, and some contemporary variant of that theme usually passes for explanation. I am always a little shocked to recall that it is only about fifty years since Paul Klee declared that it is a sin against the Creative Spirit for the artist to work when not inspired. After nearly thirty years spent in making art, in the company of other artists, I am prepared to declare that the artist has no hot-line to the infinite, and no uniquely delineated mind functions. What he does, he does with the same general-purpose equipment that everybody has, and if his use of it is in any respect unusual, that very fact points to the need for a model of image-generating behavior which concentrates specifically upon behavioral mechanisms rather than upon products.

In particular, I believe we will need to adopt a view of the artist as indulging in the generation of what I will call image-rich material as a self-satisfying procedure primarily, and only secondarily involved in the manipulation of culturally stabilized symbols: performing that secondary function, moreover, in a manner more in keeping with the essentially self-seeking character of the primary one.

You will see that I am back to the fundamental dichotomy between the internal psychic, and the external cultural determinants to an individual's behavior. The two are not available for examination in isolation of each other, for the rather obvious reason that human beings live in cultures. As far as image-generating behavior is concerned, however, it seems reasonable to speculate that image-rich material arises from the innately human domain, for the reason that the cultural determinants which act upon the individual tend, by definition, towards conventionalizing; towards the rigid binding of symbol to stabilized meaning. To reflect broader experience more accurately, we have to look, not for symbols which are unambiguously understood within their own culture - however powerfully they may function there - but for material which can flow between cultures, and which is constantly re-used to mesh with new and diverse meanings as it does so. What we think of as our culture is no more than a moment in time, a cross-section of a continuum. All but an infinitesimally small part of all the symbols and symbol-potent material which reaches us comes to us from other points in time and from other more or less remote cultural states.

In some cases, what we find ourselves responding to comes from cultures so remote that we simply have to acknowledge that we cannot possibly know what its original significance was. I am thinking particularly of the petroglyphs which are to be found throughout Nevada and California (7). We know nothing to speak of concerning the people who made them or what they made them for, or even how long ago they were made. We cannot seriously pretend even to misunderstand their original significance, and what speculation exists is based upon evidence quite extrinsic to the marks themselves. Yet the generations of anthropologists who have added their speculations to an increasing but unrevealing literature bear witness to the power of the glyphs: the power, not to communicate explicit meanings within the culture within which they arose, but to trigger and direct our own innate propensities for attaching significance to events.

To account for the pressure which these marks are capable of exerting over so total a cultural void, would we not



fig 7

have to assume that their power derives from the essentially human determinants to their making? that it reflects patterns of behavior so deep-rooted in the human organism as to be considered as constant for all human beings regardless of their particular patterns of acculturation?

The question would be entirely speculative, not to say gratuitous, if we could proceed only by the analysis of existing examples, for the reason that what is present for analysis is the object, not the behavior which generated it. Any plausible conclusion would be exactly as good as any other plausible conclusion in the absence of any possible verification.

I will not claim that my own work offers definitive verification of any conclusion, but I will claim it as an attempt, at least, to deal with behavior rather than with

objects. Analysis of a range of objects, from the Californian petroglyphs at one extreme to my own drawing at the other, has served mainly to suggest a sort of minimum configuration of deep-level behavioral mechanisms, which have then been used as the basis of a computer program capable of generating and dealing with graphic material.

In other words, the choice of mechanisms was largely intuitive and arbitrary. I suspected that I would be on reasonably safe ground if I limited myself, at the outset at least, to what I assumed to be perceptual primitives, and I selected three: the ability to differentiate between figure and ground, to differentiate between open forms and closed forms, and to differentiate between insiderness and outside-ness.

Since the choice was arbitrary, I did not think it needed further justification at that stage of a questionable undertaking. Yet I thought that actually justification could be found: both in the fact that young children evidently differentiate between closed forms like circles and triangles, and open forms like crosses, well before they are able to differentiate between circles and triangles: and also because of the persistence, throughout the long human history of mark-making, of motifs like mazes. It seemed to me that much of what we grace with the name "primitive" actually demonstrates a sophisticated awareness of the nature of the perceptual open/closed duality, for the fascination of the maze - the image, I mean, rather than the physical maze - must surely rest on the difficulty of knowing at a single glance whether it is open or closed.

The point of the strategy - the building of a computer program - was not to see whether the presence of these behavioral primitives would add a sense of authenticity to the output. It was to see whether the program could generate image-rich material in a controlled context where it would be clear that the effect was not the result of something else. That would certainly not have been the case if I had tried to limit myself to any particular set of behavioral primitives, and I have taken some care to see that I do not influence the running of the program. As it has been designed, it operates without any human assistance

or intervention. There is no way to interfere with it while it is running, and no convenient way to change its parameters before the start of any drawing.

Much more important, it has no data at its disposal: no lexicon of previously-described forms which it could pull out, run through a variety of transformations, and assemble into a picture. As a matter of fact, it has no transformations available to it, either.

An argument could be made, of course, that the whole program constitutes a process description of its output, although it would then have to be seen as the description common to an endless array of different drawings, since the program never produces the same drawing twice. But the significance of the lack of data is a more complex one. There is no difficulty about writing a computer program which generates drawings endlessly, depending at least upon what "different" is understood to mean. Here the question was whether the degree and kind of differentness would correspond to the variety we might expect from a human





Fig 9

image-maker. Would the individual drawings, generated in the absence of any knowledge of the world and its objects, nevertheless function as though they were made by a human image-maker, in the sense that they might appear to be making reference to the world and its objects?

The answer seems to be affirmative, at least to the degree that most people evidently have some difficulty in believing that the drawings (8-10) were not made by a human artist: an artist, moreover, with a distinct sense of humor and a marked tendency towards narrative.

As the prime mover of these drawings - I still have some difficulty regarding myself as their maker in any conventional sense - I find myself in a curious position involving a not-too-serious parody on the notion of divine inspiration. It takes about two weeks after seeing one of the drawings for the first time for me to lose my awareness of it as machine output. I can hardly regard it as my own, because I have no recollection of having participated physically in its making, and it seems to have come to me from



fig 10

another time and place. We might see this as a comment on the persistence of myths, perhaps. But if romance dies hard, the facts are left to be accounted for. If we find elements in these drawings reminiscent of African masks and comets, figures suggestive of turtles and submarines (10), it is a fact that the elements and figures which evoke those objects were made by the program. It is also a fact that the program knows nothing of African masks, comets, turtles or submarines.

Explaining how these effects come about in the absence of any specific intentionality is difficult, primarily because they cannot be identified with the action of individual parts of the program. There is, I mean, nothing like a submarine subroutine. In form, the program is a production

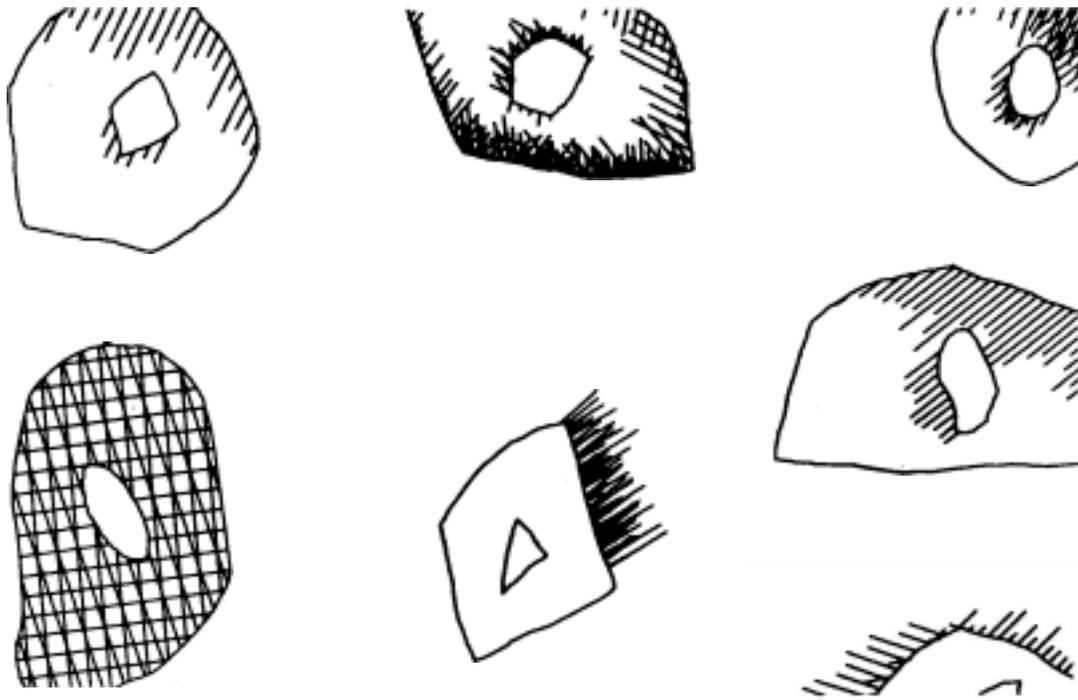


fig 11

system; and, like other such systems, this one accomplishes two things. It describes the conditions which may arise in the world of the program – in this case the developing drawing – and it lists the acceptable responses to particular combinations of these conditions.

The left part of a production tests for the patterns, the particular combinations of conditions which characterize the state of the world at any moment. The right part of a production changes the state of the world, since all the acceptable responses act upon the world directly or indirectly. The new combinations of conditions will then be trapped by other productions; and the process continues, in this event-driven fashion, from the initial empty state until one of several world states elicits the response that the drawing is done.

The left part of a production is able to recognize that a form is closed rather than open, just as the right part is able to produce a closed form, or effect closure upon an open one. A complete production might recognize that part

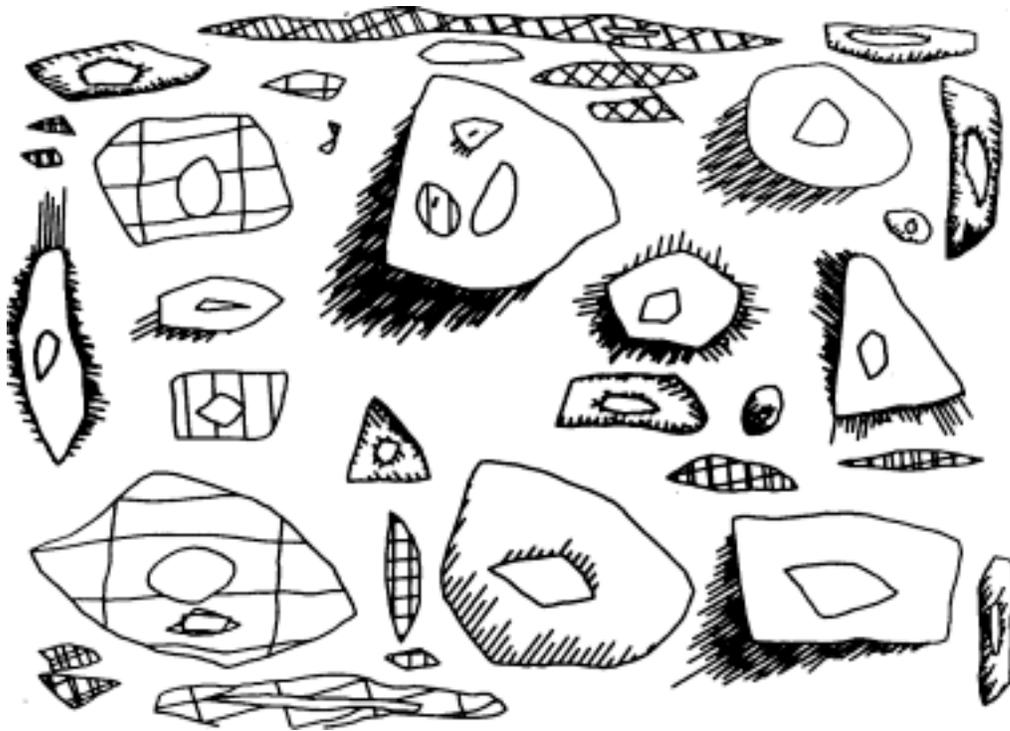


fig 12

of the field of the drawing is occupied by a closed form with another closed form inside it; and that it is surrounded by similar closed forms, all of which have been shaded in one way or another (11). And it might respond – for example – by shading the figure, leaving the inner one as a hole in the middle.

But references to closure, to space-filling, and to repetition occur throughout the production system in both the left and the right parts. They constitute, not a set of rules so much as a set of protocols, the complex intertwining of which gives the entire program its particular identity. They are best considered as characterizing the program's world rather than as controlling how the program is to behave within that world; as characterizing – if I risk anthropomorphizing a little too far – what the program understands its world to be like.

Space-filling and repetition are two of several protocols which have been added to the program since the outset, most of them simply extending upon the initial ones. I mean that

shading is a way of underlining the closedness of a closed figure, and the program now knows a number of ways in which that can be done (12). A recent extension to the figure-ground protocol requires the program to respect the territorial integrity of previously drawn figures. This one results in some of the more unpredictable and evocative configurations; though it is never easy, even watching the drawings being done, to keep track of what is causing what.

Adding a single new protocol to a program is more like adding a whole new conceptual complex to a human's world model than it is like adding a new behavioral rule, and it should not be surprising that the complexity of the drawings increases rapidly for each added protocol. This seems to suggest that the program structure is appropriate to the requirement of variety which I noted earlier, since it seems unlikely that human output increases in variety only at the cost of extremely large rule-sets.

I have not yet had sufficient time working with a reasonably well-developed program to reach detailed conclusions on the nature of that variety, and on how the enmeshing of the different protocols produces it. But it does seem clear that it is the enmeshing, not the individual protocols, which is responsible. Note, for example, that although one drawing may exhibit more sophisticated space filling - shading - abilities than another, it will not have the same evocative force as a "simpler" drawing which exercises both open and closed protocols (cf 10,12). In fact, I think there is evidence to suggest that in the presence of closed forms, open forms take on a distinctly differentiated function, providing a kind of semantic connective tissue for the semantically dominant - more obviously object-like - closed forms. It is certainly the case that the spatial relatedness of the figures significantly affects their individual reading.

There is one further aspect to the program, having to do with task-oriented behavior rather than with perceptual behavior, which I should touch on briefly. It controls the way in which the program goes about the actual production, and the physical articulation, of the simulated freehand line from which the drawings are built.

I quickly came to the general conclusion, when I first became involved in computing, that human drawings are potentially interesting to human beings at least in large part because they have been made by other human beings; and that for a machine to inspire a similar kind of interest in its products it would have to make its drawings in the same sort of way that humans produce theirs. Of course, everything I have been talking about has been an effort to elucidate what that "same sort of way" might be, but I am thinking now specifically about the lowest-level business of driving a pencil from one place to another.

What seemed certain to me, and still does, is that freehand drawing involves an elaborate feedback mechanism, a continuous matching of current state against desired end state and a continuous correction of deviation, essentially like the mechanisms we use to thread a needle, or drink a glass of water, or drive a car. Most of the time the feedback is required -- and the artist can claim no exemptions in this regard -- by the unpredictability of the equipment we use, whether that unpredictability is caused by arthritis or worn bearings, lack of muscular coordination or sloppy steering. We do not optimize in freehand drawing, and it never seemed to me that the dynamic qualities of drawing would be captured by spline interpolations. Indeed, it never seemed to me that those qualities would be reproducible by trying to mimic appearance at all.

Imagine the problem of driving your car off a main road, where you are facing in one direction, into a narrow driveway at an arbitrary angle to it. Unless you would proceed by planning your whole course in advance and then closing your eyes and stepping on the gas, you will probably be doing very much what the program does. Given the task of getting from one place, facing in one direction, to another place and facing in another direction, it never knows how

to accomplish the entire task, but "imagines" a series of temporary destinations, each of which will bring it a little closer to approaching its goal from the specified direction (13)' A degree of randomizing is provided as an

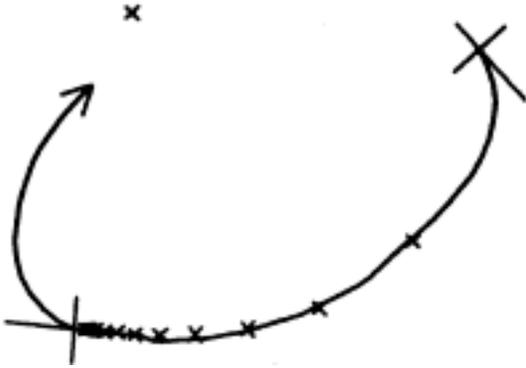


fig 13

analogue for arthritic joints, and as it never had any precisely defined path to follow anyway it corrects for accumulating discrepancies only when they become big enough to jeopardize its chances of ever reaching its final destination.

It never knows in advance what will constitute a complete path, and it never fails to complete its path. This part of the program is non-trivial, and certainly not optimal, involving as it does a complex series of decisions for every one of the small line segments which go into the building of a line. But I believe the simulation is a good one, and I have found it possible, moreover, to modify the character of the line - the artist's "handwriting" - by the manipulation of such thoroughly practical factors as the rate at which sampling is done, the suddenness with which correction is applied, and the frequency with which the program sets up new "imagined" destinations along its path.

It seems to me that most of the things one might say about image-building might be said equally about image-reading. The reason for this, I think, is that the element common to both -- the propensity for attaching significance to events, for endowing entities with identities -- is also an overwhelmingly important one. It is not the unique property of artists, obviously.

This is not to say that the identity which the viewer attaches to a complex of marks is exclusively a function of the viewer's propensity, or even that any complex of marks would serve equally well to trigger that propensity. The natural world is full of complex forms, and if we sometimes play with them -- clouds, for example -- we are well aware that their "meanings" are our own invention. Marks which we recognize as being man-made, on the other hand, -- and in particular those man-made marks which we see as arising from an intent on the part of the maker to communicate -- these we treat in a special way, not merely assigning significance to them but insisting that that significance has been carried by the marks themselves.

I believe that in searching images for evidence of their origins the mind is surprisingly literalistic. If a machine program is able to produce image-rich material, it does so by virtue of persuading the viewer that the maker was a human being living in a human world, and that his intent was to communicate something about that world. The assumption of intentionality precedes the "reconstruction" of intent.

In this case the simulated perceptual mechanisms give evidence of the underlying humanness of the drawing's manufacture and the drawer's world -- though perhaps any other set of reasonably low-level mechanisms would have served equally well -- and the constant complex decision-making which actually takes place, and which is clearly evident in the articulation of the line, confirms the viewer's belief in the artist's intentionality.

This conclusion is not adequate to account for the more highly particularized readings which seem to attach to the drawings -- notably the humor and the sense of narrative --

and I do not know at this stage how they are to be accounted for.

It became evident from the questions and the private discussions which followed this paper that my use of the label "protocol" had done more to confuse than to elucidate the conceptual unit to which I had applied it. Reviewing that usage, it becomes apparent that my understanding of what the program is doing – what roles the different elements in its structure play – has shifted with time, and I have been careless enough to carry over to a slowly emerging construct a term inappropriate to it, but unfortunately still more or less appropriate to something else. The underlying confusion has been my own, of course, and I am glad to have been presented with this opportunity to try to resolve it.

Given the choice between rewriting the paper and extending it with a post-scripted commentary, I have chosen the latter course. This gives me the chance also to deal in a more measured fashion with one particular question which is evidently quite troubling to a good many people.

I suggest above that what I call a protocol is best regarded as characterizing what the program understands its world to be like, not as a rule which controls how it is to behave in that world. A rule is expressed within the program by a production. A protocol is not fully expressed by a production. I would not want to change any of this - except the use of the word "protocol" itself - but the problem is that nothing has been said about the structure, or what we might call the dynamics, of the characterization. In the absence of any overview clearly differentiated from the rule-oriented schema to which the characterization must obviously relate, the mere assertion that a protocol is not simply a rule is hardly sufficient to expunge the sense that it is.

So be it: let me return "protocol" to the rule-oriented domain whence it came. In its place, and hopefully more fully expressive of the conceptual complex it is meant to carry, I will use the term "epimorph".

An epimorph characterizes what the program understands its world to be like, and the machine draws in a human, or quasi-human, fashion because its set of epimorphs are closely modeled on human epimorphs. We might go as far as to say that it exercises a subset of human epimorphs. In dealing with the dynamics of the characterization process, then - and thus in attempting to elucidate what an epimorph is - it may prove more revealing to consider an example of human, rather than machine, performance. Here is one taken from the drawing class mentioned earlier.

A brief background account is in order.

Two weeks into a deliberately dislocative class - people bring such rigidly formulated notions about drawing to a beginning class! - one struggling student volunteered the view that drawing was, as far as he could tell, "just a question of getting from one point to another". Always happy to take what is offered, I proposed that in that case they might get into the business of drawing more freely if they didn't have to worry about the points. Each of them could provide an array of dots for someone else, who would then only have to figure out how to get from one to

another.

In practice, it required fairly rigorous measures to ensure that these dot arrays did not carry any representational weight of their own to constrain subsequent performance. Eventually we had two sets of drawings, thirty-four in all, pinned up for examination, and before any discussion began I asked the students whether they could write down the rules which they had followed in joining up the dots. They all wrote down the same three rules! - 1. see if you can see an image in the dots, and if so draw a line around it: 2. if you can't see an image, draw closed figures anyway: and, 3. if you can't do 1 or 2, fake it. "Faking it", on questioning, turned out to mean using open structures like short straight lines, zigzags, and so on, as space filling.

Examination of the drawings themselves showed that there were several other rules of a more surprising kind operating. Consider that any dot in an array might potentially become the junction of an indeterminate number of lines joining it to any number of other dots. Of the simpler cases, the order-two case denotes a dot on a continuous line, the order-one case marks the end of a line, and the null case is a dot which has not been joined up to anything. Karin's "eyes" would be an example of the null case.

Since the drawings all contained between a hundred and two hundred dots, we might guess that there would be considerable variety in the numbers of lines joining at these junctions: in fact, we found only three cases of order-four, and only two cases of order-more-than-four, junctions, in the entire set of thirty-four drawings. Over 99.5 % of the dots had three lines or less attached to them! (A similar situation will be observed in the drawings of both Karin and Sherry, figures 3 and 4.)

The students were certainly unaware, until it was pointed out to them, that their behavior had been constrained in this way, and were even a little resentful of the suggestion that they had done anything according to rules of any sort. Yet, curiously enough, there were a few cases where "extra" dots had occurred when two lines had been allowed

to cross, and in all these cases the students concerned reported a strong sense of having done something wrong, broken some powerful though unstated rule. The class as a whole evidently recognized an unstated interdiction against crossing lines also, and unanimously agreed that these "extra" dots should not be counted as order-four junctions.

Consistent though this behavioral pattern was, it only required attention to be focused upon it for it to change. The discussion which followed the making of these drawings evidently identified "junctionality" as an issue, and although nothing was said about what might constitute acceptable behavior in relation to this issue, the drawings which followed in subsequent weeks all contained a much richer distribution of order-more-than-four junctions; we discovered also that their use involved increasingly complex, but hardly less consistent, rule-sets than we had found at the beginning.

We need not go into detail here on the precise nature of these new rules. The point is that they could all be described by a production-like paradigm involving assessment of the current state of the drawing - in relation to junctionality among other things ~ on the left side, and some action resulting in a change of state - through the manipulation of junctionality among other things - on the right side. The notion of junctionality itself would not be adequately expressed by any one of these productions, however, and it clearly exists on a "higher" hierarchical level than that of the individual productions. It has become one of the issues which the student believes to be significant in relation to the domain of drawing, and thus characterizes what he believes that domain to be like. It is in this sense isomorphic with those other issues of territoriality, openness/closedness, containment and repetition, which I said characterized what the program understood its world to be like. It will be clear from this account that of these, at least openness/closedness is also an active epimorph for the human: but I am sure that more extensive evidence will be found in a wide variety of material, and in domains not limited to drawing activity.

One of the questions I was asked -- not for the first time, by any means -- was: am I proposing that the machine program constitutes a model of human creative behavior? Is it a sort of automated surrogate Harold Cohen?

A full answer would go far beyond my present scope -- and, indeed, my present abilities -- and would involve all those other troubling philosophical questions which the existence of the computer inevitably raises. A short answer would be that human beings live in a real world, and their internal representations of that world include reference to its objects: the current state of the program knows nothing of the real world or of its objects. Human beings learn from experience: the program begins each new drawing without any memory of previous drawings, and with its production system unmodified by having made them. In these and in other respects the machine's performance is not merely less than, but is unlike, human performance. It should be stressed, however, that these are limitations in the current state of this program, and are not to be regarded as intrinsic to programs in general.

Most searching questions about the nature of the machine turn out to be questions about the nature of people, and this one is no exception. Before we could venture a more complete answer we would need to consider what we really mean by creative behavior, for if that is to be judged exclusively in terms of the manifest results of its exercise -- we know so-and-so is creative because he makes a great many original images -- then clearly the machine is extremely creative. Its drawings are probably as good, as original, as any I ever made myself, and I am hopelessly outclassed by it in terms of productivity.

But once we have stripped off these layers of the artist's activity which have to do with marketable objects, with the desire for approval, for fame or for notoriety, with propaganda for this religious belief or that economic system: once, in short, we have stripped off the artist's public and cultural functions, how will we characterize the remaining private, essentially self-serving, functions? What does the artist make images for?

My own view can be stated briefly and without oversimplifying too far. I believe that the artist is engaged, as everybody else is, in building internal representations of his world, and that his behavior is remarkable in only two major respects. The first – and this seems to me to be a feature common to art-making, science-doing, philosophy, mathematics, and most other higher intellectual pursuits – is that the formulating and continuous reformulating of mental models is carried on as a foreground, and as a highly structured, activity: not as a background activity. The second is that he exhibits a high level of preoccupation with the structure of representation as such.

In neither of these respects does he require special mental equipment, and indeed I would assume that the cultural value of his activity, the extraordinary regard in which images are held, rests upon the fundamental normality of the mental functions exercised. I mean that the basic structure of all internal model building is the assignment of associative reference: what we might call the "standing-for-ness" principle. We would not be going too far to regard art as an endless explorative game built around the presumably universally human fact that things can stand for other things.

The playing out of this game produces images, normally embodied in objects, which may be valued by the culture for any of a number of reasons. For the artist, it is the playing out of the game, and thus the making of the object – rather than the object itself – which is important. If object-making is the means to an end, the end is not the object – art objects are interesting to the degree that they stand for something outside themselves – but the continuous development of new moves in the game. Externalization is a part of the artist's methodology in the building of internal representations of his world: a world which includes representations as a central feature.

We are now in a position to generate a slightly more complete answer to the original question, and I think we will find that the view which the question proposed – that the program is an artificial artist capable of creative behavior – is both more than and less than adequate. The

program does not develop new game-states: it plays the legal moves in the current game. It says "Let me tell you about my world", but rich though that world may be, the telling does not result in any further enrichment. We thus have no reason to say that the machine has any interest in the one feature I have chosen to regard as fundamental to human art-making – the continuous development of the internal representation of the world.

To this degree, it is clearly an inadequate model of human performance: which is not to say that no program could ever provide an adequate one. On the other hand, it does not merely model the playing of legal moves in the game, it actually plays them. To this degree the program is not a model of human performance at all. It carries out a real, and rather extensive, part of the art-making procedure, and its output is in every important respect interchangeable – both culturally and privately – with output which might result from more orthodox art-making procedures. MY world changes as a result of the program telling about it, and in the long term the program changes also. I assume from this that I will go on working on the one program indefinitely, without ever feeling the need to abandon it and start on a completely new one.

Some caution is in order. I have reached this point in many conversations to be told "Oh, you mean that the computer is just a tool." The answer to this is that the advent of the electronic computer requires a total rethinking of what tools might be, for if the thermostat and the speed governor are exactly equivalent to biological feedback systems, computer programs are potentially exactly equivalent to intellectual feedback systems. We have a long way to go before we fully comprehend the shift in significance of "tools" capable of the independent exercise of reason.

I have said several times that the limitations attaching to this program should not be regarded as fundamental limitations in programs. I do not know what will change, for example, or how they will change, when this program does have some knowledge of the world, and can make decisions about the drawing in terms of that knowledge: or when it can use its memory of past drawings as a determinant in

building new ones.

Prediction is a hazardous game, and I will limit myself here to only one. I do not believe that any program will ever produce art unless it was written by an artist – as the words have been defined by this discussion – and its running serves a vital role for that individual in the changing patterns of his internal model building. The Sci-Fi fantasy of putting an artist's "genius" on tape and flooding the world with his work after his death, or of becoming a great composer in the twentieth century by writing a program to generate Bach: these merely reflect the confusion of art with its objects.

