Organism, Environment, and Intelligence as a System

JOHN PLATT
Mental Health Research Institute
University of Michigan, Ann Arbor

Someone has said that biology can be exhaustively subsumed under three headings: evolution, morphogenesis, and thought. Evolution and morphogenesis have already been discussed and so it is a pleasure to turn to the subject of thought. What I propose to do is tell about some findings, attitudes, and insights that come out of work in our laboratory. These are largely theoretical, and they largely deal with models, but it seems to me they have considerable application to the relation of an organism, or more exactly of "a sensory-motor decision-system," to its environment.

LIMITATIONS OF AMPLIFIER-INPUTS

The first aspect of such a system which I want to emphasize is that inputs into the sensory network involve *amplification* and *selection*. For example, the visual receptor cells, the rods and cones of the eye, are sensory amplifiers. It is now believed that the rods probably operate at the "one-quantum level." This means that a single quantum of light coming into a rod can be detected, producing an amplification, which is of the order of 10^7 times (in electrical voltage), in the first nerve impulse that goes out into the retinal network and then is transmitted down the optic nerve. It is interesting to make a comparison between this amplifier and a little electric photomultiplier pulse amplifier. The rods and cones are very small, only some 50 microns long and 1 to 3 microns in diameter, and this factor of 10^7 in amplification in such a small system is a very impressive performance.

But these amplifiers, like all other sensory and electronic amplifiers, have necessary limitations on what they can detect about the external world. The limitations come from the fact that they are *macroscopic*, that they have a *time constant*, and that they have some *noise* background so they are not absolutely reliable. I would also emphasize that they are *selective,* and that they are *unidirectional* in their operation. Input into the photocathode or a phototransducer molecule (rhodopsin) is amplified *in one direction* down the system, and then produces an output which comes out at the other end. This chain cannot be reversed by putting in a light quantum at the output end and expecting to get an amplified signal at the input. I suspect that some of the aspects and limitations of perception which are important to us may depend on such limitations in these elementary amplifiers which mediate between us and the environment.

SENSORY-MOTOR DECISION-NETWORKS

Secondly, I want to emphasize some of the new results on the combination of input-amplifiers and output amplifiers into sensory-motor decision-networks, and on learning networks in particular.

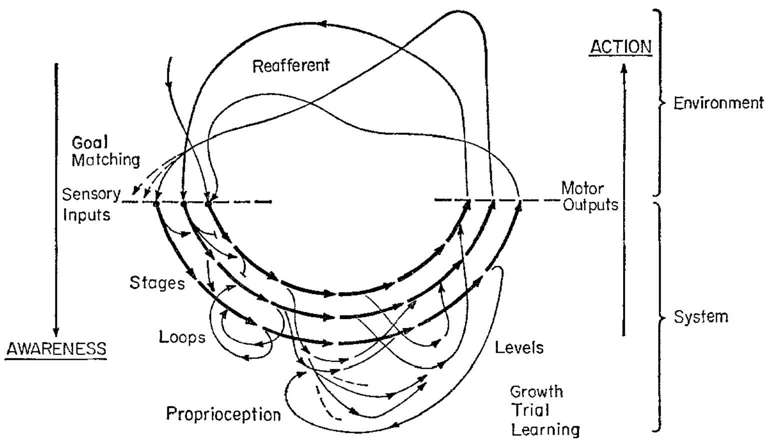

Fig. 1. A sensory-motor decision-network interacting with the environment.

Figure 1 is a schematic diagram of an organism detecting its external world. It is an attempt to diagram some of the essential relations for an organism as we now know them. In fact, what I have in mind is specifically an artificial organism, or "automaton," which is capable of surviving in certain types of environments. I am interested in the parallel between such an automaton and a small elementary nervous network which has certain properties such as learning processes.

One should not belittle the value of such diagrammatic representations. The specification of models of this sort has played a great role throughout history in providing insights into the nature of man, from the Cartesian automatons to Wiener's feedback models. When one wants to study the brain, for example, a good model can be illuminating, not because it is exactly like the brain but because it suggests the kinds of questions to ask and the kinds of relations and details to look for. Wiener's cybernetics, or feedback model, to my mind is one of the two or three most important ideas of the twentieth century in terms of its implications for all sorts of fields. What we have to have now is a model which includes not only the essential elements of these useful models of the past, but some important new elements as well.

In particular, this diagram is intended to represent a *parallel-processing cybernetic system.* All the millions of sensory input cells are symbolized by three little circles on the left side of the diagram which are on the boundary between the organism and the external world. From these sensory cells, in this schematic model, there emerge first-stage neurons (below left). These may be connected to second-stage and third-stage neurons, and so on, and there may be cross-connections as shown at the bottom of the diagram. There may be other features, such as proprioceptive feedback loops, and there may be some kind of growth of new connections (shown at the bottom) as one goes to higher-order organization of events.

All these neuron chains lead eventually to motor outputs (right side of Fig. 1). In fact, unless we do micro-electrode studies, essentially all we can learn about the internal workings of such a decision system is in terms of its motor outputs. These motor outputs feed back into the environment and lead to various environmental consequences (top of the diagram) which may or may not feed back into input in such a way as to modify this input.

REAFFERENT STIMULATION

The reason for emphasizing this particular parallel-processing network model, in which the inputs lead to motor outputs, which lead back to inputs, is because it is a good representation of the important phenomenon of "reafferent stimulation." I think it was von Holtz who first suggested the term, but recently the work of Richard Held and his colleagues in the

Psychology Department of M.I.T. has shown the great importance of this process. What reafferent stimulation means is a change of the input field as a result of the motor output of the organism or animal itself.

Held's work has shown that reafferent stimulation is essential to the organization or adaptation of the visual field. This had already been shown in a number of small ways. For example, Ditchburn and Riggs and their coworkers showed fifteen years ago that continual motion of the eyeballs is necessary for visual perception. They measured the motions by putting contact lenses on the eye, with a little mirror on the side of the lens to reflect a beam of light. They showed that when the subject thinks he is fixating most steadily on a point source of light, its image is in fact moving around over the fovea, which is a region about 1° across at the back of the eyeball. The movement consists of a series of drifting motions on which is superimposed a tremor motion of about a minute of arc, combined with occasional jumps or "flicks."

These motions of the image over the retina can be canceled by an ingenious mirror arrangement, so that spots you are looking at move just as fast as your eye moves. Or they can be canceled even more successfully by mounting the little spot of light, or a microscope lens and a target pattern, on the contact lens itself, so that it all moves with the eye. The lens combination has no appreciable weight, so it does not significantly affect the motion of the eye, and it is not uncomfortable physically.

The remarkable thing is that with such a "stabilized image," vision ceases. For example, if the fixed-pattern apparatus is put on in the dark and if the light is turned on, the pattern can be seen for a fraction of a second and then it fades out. Only when some motion of the image is allowed is it seen again. What this implies is that at this eyeball-level, motion is necessary for visual perception. Presumably the stabilized image fades out for physical-chemical reasons because the cells adapt, or come to a "steady state."

FUNCTIONAL GEOMETRY

I have used this principle myself in an attempt to solve what I think is a most interesting and classical visual problem: how do we see straight lines? I do not think my own solution is widely accepted in the field, but nevertheless I think it has some important properties which I would like to bring to your attention.

The problem of how we see straight lines with a highly variable cellular retina is a difficult one by all past methods of approach. I think it is almost impossible to understand if we simply think of a static image of the outside world falling onto the retinal array, because the retinal arrays are genetically different from one person to the next. But even if they were not, it would be quite surprising if during childhood development, every cell came to lie in exactly the same micro-relation to the other cells, or if any micro-relation like this should be individually predictable to a high enough accuracy. It is worth remembering that the accuracy of our perception of straightness, at least in detecting the lateral displacement of a line segment, is of the order of 2 seconds of arc—much smaller than the diameter of a retinal cell. Certainly when we look at the retina through a microscope, the cell positions look very irregular.

It therefore seems difficult to believe that a *static* map on the retina could be interpreted by you and by me to indicate the same object pattern, such as straightness, because the image will fall on one set of cells in your eye and a different set of cells, differently arranged, in my eye.

But if we use reafferent stimulation—that is, a *dynamic* mechanism in which we move the images back and forth on the retina—the problem suddenly becomes much more simple. For example, a way in which we can tell whether the top edge of a blackboard is straight is to look at the left side of the edge and then to look at the right side. If the image of the left side falls on a certain set of cells in your eye, and if it continues to fall on the same set of cells as you move your eye across to the right side, then the line is straight. For no S-curve or general irregular curve can continue to fall on the same array of cells in your eye as you move it.

In short, I am suggesting that the simplest way to organize the geometrical space of the external world is by looking for *invariances under displacement.* I call this the principle of "functional geometry," meaning a dynamic approach to geometry rather than a static approach in terms of mapping. The operation of "invariance under displacement" can be used by any eye and retina to single out as unique such important geometrical pattern elements and regularities as straight lines, parallel lines, and equidistant lines, as well as things like uniform circular arcs (provided you can rotate your eyeball about the optic axis, as indeed you can).

It is interesting to see that this method of extracting pattern regularities is very similar to the methods used for generating mathematically perfect surfaces, such as spherical and plane surfaces, in high-precision optical work. John Strong, the designer of high-precision ruling engines, has described in detail the methods of grinding and lapping for generating perfect spheres, perfect planes, perfect helixes, perfect equidistant spacings of a ruled grating, and perfect gears with high-precision angular spacings of the gear teeth. Basically, his methods do not involve measurements, but simply involve sliding and lapping one surface, such as a spherical surface, over another which just fits it.

In this kind of displacement-invariance method, you do not use any rulers, and you do not have any centers! The spheres define their own centers and radii as they are polished, and the numerical values are not even measured. If you try to specify the center, it turns out that you damage the approach to sphericity because you are "over-determining" the system. It has to determine itself. In commenting on this elimination of measurement and external centers, Strong says, "the methods of highest precision are all primitive methods."

Does it not seem likely that a biological system would need to use some primitive general high-precision method of this sort? For the eye has no x-axes and y-axes for Cartesian-coordinate determination of the patterns of figures. It has no method, as far as I know, for determining the "shortest distance between two points."

In fact, the elimination of these classical bases of measurement and construction means that functional geometry gives us a new axiomatic basis for geometry. We suddenly realize that Euclidean geometry is "string geometry;" it is the kind of geometry you get by drawing on the sand using a piece of string. This can be seen in the definitions: a straight line is the "shortest distance between two points"—of the kind that is represented by a stretched string. A circle is the "locus of points equidistant from a given point"—using a length of string as a measure of distance.

Cartesian geometry, on the other hand, is "box geometry." It is the geometry of a European house, in which the walls are made of bricks that step up 1, 2, 3, 4, 5 . . . units. The floor is made of tiles that march out 1, 2, 3, 4, 5 . . . units, north and east. (I suspect that the Greeks, with the bowed-out columns and arched floors of their temples, could never have invented Cartesian geometry.) But now we come to functional geometry, which differs from these static, locus-of-a-point, geometries, and which has no axes, no centers, or "origin," no measures of distance, no guarantee that some primordial axes are at right angles to each other— and which gets along extremely well without all this initial apparatus.

Incidentally, this even throws a new light on the "parallel postulate." In functional geometry, the parallel postulate is immediately provable, because if two lines are invariant under displacement along them—that is, if parallelism is defined by that kind of invariance—then if they do not meet in one region, they will never meet, because they are invariant. Therefore two parallel lines so defined will never meet. What I suspect is that Euclid or his predecessors had a kind of intuitive biological sense of this sort. Perhaps Euclid felt that two parallel lines ought never to meet, and since this cannot be proved by any method such as the "interior-angle" construction from the other axioms, he felt it had to be added as an additional axiom.

The kind of discriminations that can be made by functional geometry raises an interesting question. The distinction between a "straight-line" displacement (x- and y-axes of the eyeball) without rotation (about the z-axis), and a "slightly-curved line" displacement with a little bit of rotation (about the z-axis) does not depend on functional geometry, since both patterns may be equally invariant under the right muscle-motions. Rather, it depends on the accuracy of the "analog system" monitoring muscle movements, which tells you how much the muscles (of the z-axis) are moving. We thus have a lower-accuracy analog system for detecting some features of the pattern and a high-accuracy invariant array system for detecting other features of the pattern. This is like the boy polishing a spherical mirror, who cannot control the curvature very accurately even though he gets a very accurate spherical surface in every case. We cannot in fact tell the difference between a straight line and a slightly curved one unless we have another line nearby to compare it to (which converts the judgment from an analog judgment to an invariance-judgment). The analog and the invariance mechanisms are then two mechanisms working on top of each other. The activities of the muscles of the eye (analog) do not in themselves define a very accurate axis system for the movement of the eye, and when you are in a darkened room the image seems to drift around in quite an arbitrary way. But when you look at a circle or an object with geometrical regularities, the object itself, by its invariances, can determine its own centers and axes with high precision.

PROPERTIES OF INVARIANT-SEARCHING SYSTEMS

I want to emphasize now some properties of this kind of system. If it should be used in perception, it would have several rather fascinating characteristics. For one thing, the perception of regularities by this method does not depend on the location of individual cells on the retina. All you ask is that the array of inputs is the same before and after the displacement-motion. Also, the accuracy of the perception does not depend on whether the image is straight or distorted—because if the image of the left upper edge of the blackboard falls on an S-curve on your retina when you are looking at it over here, and then the right upper edge falls on the same S-curve when you look over there, then the line is straight! The question is simply invariance, and not the shape of the image.

And just as it does not depend on the shape of the image on the retina, so it does not depend on the shape of the image farther back in the brain, on the cortex. The nervous network carries back a kind of distorted retinal map to the cortex, but we know that it is split apart onto the right and left half of the visual cortex. But we see that this split in space is irrelevant, as long as we have an invariance in the network under the scanning motion.

The result is that with this kind of mechanism straightness is a property in the *external* field, not in the internal field. It is a property in a *public space*, in spite of our private differences in detailed anatomy. Held has therefore described functional geometry as a method which "transcends the differences in individual anatomy." This transcendence of anatomy is obviously extremely important if we are going to be able to have an ostensive public language. Only so can there be public reference, in which *I* can point to something, and have some assurance that *you*—or a young baby—may look in the same direction and may see a regularity similar to mine, in spite of the fact that your network and his network is different, each of our arrays of cells is different, and so on. Only by some such physical transcendence-method can we arrive at interpersonal regularities in an external public space, which can then be a basis for ostensive definition and for language.

INTERACTION WITH THE EXTERNAL WORLD

There are further interesting properties which appear when we consider how such a system interacting with its external world must work. Referring to Fig. 1, it can be seen that in this kind of relationship the external world completes the feedback loop. That is to say, we have half of the feedback loop inside the organism, while the other half is outside the organism. The result is that if we look at an object whose input-signals have just begun to come in, we do not see the object unless the eyeballs move, unless our motor system is manipulating the images across the visual field. In this sense, we do not see "an object"; we only see a family of invariance-relationships which are determined in part by self-motion.

We might say with Martin Buber, then, that our perception depends on our having an "I-it" relationship to the object. Or to put it another way, we can organize the external world *only by manipulating it,* only by putting out our own amplified motor outputs in such a way as to alter what is in the field. It would not be wrong to say that in this kind of system awareness depends on action. There is no such thing as absolutely passive observation.

Another important aspect of a parallel-processing cybernetic system like this is that it is an amplifying system. I have spoken of the amplification at the inputs, but we also have a further amplification occurring in the motor-operations of the outputs. These amplifications modify the external world; for they are designed or programmed to do that, and it is necessary to do that for the survival of the organism. Such an amplifying organism is a center of change of the environment, creating changes which can be amplified up to the maximum power that the organism can control.

One new conjectural theorem for such complex systems is that it is probably not possible to predict the behavior of a given decision system by another decision system of the same order of complexity. The reason is that the second decision system cannot determine the initial state of the first decision system, and cannot know what the detailed quantum inputs are into this first system. This is because, in the first place, the inputs may be private: they operate at the one-quantum level, and one decision system may respond to one or a few quanta that signal, say, the presence of a red light, while there is no reason why the other decision system, in a different region of space, necessarily sees these quanta at all.

Similarly, the internal structure of a decision-network is private, for it is very much more complex than the number of sensory detector cells that a similar network has to look at it with. Thus, in the case of human beings the number of input cells is of the order of 10^8, while the total number of neurons in the human decision system is of the order of 10^{11}. And the total number of synaptic connections between the neurons in the brain is of the order of 1000 times more numerous—since the average cell is estimated to have on the order of 1000 synaptic buttons—so that you have some

10^{14} specific or not-so-specific synaptic junctions in the network. The result is that with my 10^8 little neurons of the eye, even if every one of them were detecting a different neuron or synapse in your brain, I could not even in principle detect the total initial state of your 10^{11} or 10^{14} neurons or synaptic connections at a given instant. Possibly there is no law in principle against the "mutual determinability" of two complex decision-systems (although I think there is), but there is certainly a kind of practical "forbiddenness rule" that one decision system cannot know what another decision system of the same size is thinking. Perhaps it may be possible for us with our 10^8 input-elements to see what a tiny earthworm's network is doing, but not what any network of comparable size is doing.

EPISTEMOLOGY OF DECISION-SYSTEMS

There is no such difficulty in physics and chemistry observations. The reason is that the rules and laws of physics and chemistry are "low-information laws" in which we detect the motion of a few stars or atoms or indicator-needles. With our 10^8 detectors, we can detect redundantly, and with a good deal of invariance-checking, the motion of one billiard ball. And even when we come to the complex things in physics and chemistry, such as the observation of water boiling or of compounds of various sorts, these are still low-information subjects compared to biology. And they are very low-information subjects compared to the structure of a human brain. So I think many of the things that have been said in the past about the epistemology of science are really based on what we can know about low-information physics and chemistry. The epistemology, and our ideas of what can be operationally observed or predicted, are going to have to be modified when we jump from the small thousand-element or million-element systems of physics and chemistry over to the 10^{11} and 10^{14} element-systems of the human brain. This is such an enormous jump and goes so far beyond the number of sensory elements available to study such systems that it will require a new attitude toward what we can know about each other.

Of course there may also be statistical regularities and predictabilities in such decision-systems, and invariance-regularities of the type I have described. But there may also be assassination behavior which occurs only in the distorted brain of one particular person, and which is almost unpredictable in the billion-element conditionalities of its success or failure, but whose effects may be amplified up to the total energy that the human race can manipulate.

It seems to me that this amplification in manipulating the environment leads us into a future which is branching out from each decision-system in an essentially unpredictable way. As a result, I am more sympathetic than most philosophers are to the recent attack by Milič Čapek on what he calls the "myth of frozen passage." He is attacking the traditional Minkowski-type of four-dimensional time-space diagram, in which trajectories of particles in time are represented as being essentially completely predictable from the past. All pasts and all futures are thought of as being simultaneously "present" together in this frozen space of the diagram, which is perhaps not so different from the old idea that they were simultaneously present in "the mind of God," who is thus able to know all future destiny in advance. But the present thoughts suggest a quite different formulation, one in which future trajectories, with high amplification-manipulations of energy and materials, are branching out from this present time at any point where complex decision-systems are operating, in ways which are essentially unpredictable in detail by any other similar decision-system.

A decision-system is continually modifying the environment around it. One arrives at a kind of existentialist picture of the operations of a decision-system, for which the only time is "now." As it uses its existing structure to operate on and amplify the existing and incoming environmental influences, it reshapes the world around it into a future of its own decision-determination.

A further curious point is that the time and space of such a sensory-motor decision-system in their manipulation-rules come to be rather different from the time and space of physics; in some sense they are not Euclidean.

We cannot speak accurately of a time, t, or an instantaneous time for a decision-system, because it takes 30 milliseconds or more for the nerve signal to go through me, so that my output at a "given" t is summating inputs over a range of times before, and in fact it is uncertain just where and when we should say, "*this* is the output." The time is not the time of a lightning flash, as in physics, but is rather a distribution of times-past, including all those that are stored or reverberating in the structure so as to affect the output. This is true also of space: the inputs are not inputs from a point locus; rather they are inputs from whole figures, from the whole array of inputs. And the outputs are not outputs at a point, but are outputs which fill the room, as my voice fills the room.

When is the instant of time? Where is the point of action? There is no *point* of action. What we are doing is operating on a kind of sensory input-totality and manipulating it into another kind of totality at the outputs. I am not saying that this cannot be reduced and analyzed in detail point by point; and I am not saying I do not believe in differential equations. What I am saying is that the philosophy of operation of this kind of decision-system is not the philosophy of operation of the locus of Euclidean points moving on sharp paths through a three-dimensional space continuum.

The result is that this kind of decision-system operates in two worlds, so to speak. The first is the world of external, or objective, science, of the physics of the environment around it. This is the world in which the decision-system "prepares a state," lets the external objects run on without interaction, and then sees what happens. It is the external-environment part of Fig. 1, in the upper half, only weakly interacting with the organism. The world of physics is a world of weak interaction.

But the other world, the world of the bottom half of the diagram, is not the world of physics but the world of cybernetics. It is the world of goal-directed behavior, of manipulative, purposive behavior which is trying its hardest to amplify and to interact strongly with the environment. It is the world which is not primarily concerned with scientific and detached observation and information, so to speak, except as these contribute to its personal knowledge, values, purposes, decisions, and acts. This very instrumental attitude toward science is implied, of course, in our basic science textbooks, but implied in an oddly suppressed way so that the students tend to forget their own human operational role—their decision-system role—in creating and evaluating science.

THE THREE METHODS OF PROBLEM-SOLVING

There are three methods of problem-solving which evolved in the course of evolution. The first is the method of problem-solving by survival. This is the "phylogenetic method" in which all the moths that didn't match the color of the bark got picked off by birds, and what is left is a species which has solved the problem of matching the color of the bark.

The second method is problem-solving by individual learning. This is the "ontogenetic method." Whereas species survival involves the selection of chromosomes and DNA, individual learning involves the selection or connection of learning neurons in a learning nervous system. In this case the beast does not have to fall over the cliff before he is able to draw back from it, because he can see the edge of the cliff in time, or he realizes that he has been in a situation like this before, so he pulls back.

The third method is problem-solving by anticipation. This comes only with the development of symbolic manipulation, so that by abstract analysis one knows the laws of nature which can be extrapolated into the future. This is of course the method of science.

The difference between these methods is illustrated by the first Sputnik that was sent up. It was not one of thousands of Sputniks shot into space, of which all died except one which happened to go around. Nor was it like the rat finding its way through a learning maze, with the rocket going first too high and then too low until it finds out which orbit is satisfactory. No, the rocket was programmed in advance, with problem-solving by anticipation. An on-line feedback-control program was set up, and the laws of physics were used in the computers and the feedbacks to determine into which orbit it should go; and it went in the right orbit the very first time.

The importance of these considerations is that more and more of our problems are problems that we have never met before; and they are problems that will have to be solved by anticipation by science—if they are to be solved at all before they kill us.

This anticipatory manipulation of the environment by complex decision-system intelligence is something that has come late in the history of planetary evolution. From a broader point of view, Lederberg has described three great stages of evolution, starting with the preorganic molecules. He calls them stages of "chemogeny," or build-up of complex molecules in the sun's radiation-field; "biogeny," after self-catalysis begins, with natural selection; and "cognogeny." What I have been describing here is the development of systems making a collective transition over to the cognogenic stage of evolution, in which we are doing abstraction of natural laws and symbolic manipulation with enormous power control over the environment. Eventually, of course, cognogeny will include the extensive development of artificial sensory-motor decision-systems, or automata.

It is philosophically interesting to inquire whether there may be some kind of convergent evolution—in different solar systems, for example—to the cognogenic stage. One sees the convergence of wings in reptiles, birds, and bats to structures built in different ways but having similar longitudinal struts;

surfaces that have a "camber," a similar length-to-chord ratio for animals of a given weight floating in our particular density of air, and so on. There is a convergence to certain inherent or predetermined regularities of wing formation which are necessary for aerodynamic success.

One can say the same things about the convergence of eye-forms, which have been invented independently seventeen times in evolution, according to one list. Again they converge to a small number of forms, with all of them in the higher organisms, I believe, having some close-packed directional-mosaic of cells like a retina. We suppose that this happens because this is probably the best way, if not the only way, to see simultaneous external relations in many directions—by using a close-packed parallel-processing input array. The physics of the situation, the environmental necessities, channel what is successful into certain directions.

Religious people have interpreted this regularity of efficient form as an "argument from design." And I think a scientist could even agree that the wing of an animal of a certain size is "predestined," at least in a metaphorical or phsyical-necessity sense, not to look like a club. It is predestined to have certain aerodynamic parameters.

I am emphasizing this because I want to go on to ask the question, Is intelligence, and the emergence of the cognogenic stage, also a case of evolution toward a "predestined" set of parameters in the same fashion? That is to say, once complex decision-system organisms have passed a certain threshold in manipulating the environment—which might be a threshold of the use of tools or speech or fire—isn't it almost "necessary," in this same sense, that they go on to develop a manipulating environment-organizing symbol-using brain of the type we have? Could one talk about the possibility of convergent evolution in this sense toward intelligence? I think this is not an absurd or meaningless question, but one which could be solved by theory and experiment. And I think it is a question which biologists have an obligation not to deny or erase, because it is a question that inevitably comes up in discussing evolution, in discussing the meaning of the appearance of man and of intelligence with citizens and religious people and philosophers.

COLLECTIVE INTELLIGENCE AND THE COLLECTIVE FUTURE

We are on the threshold of a very peculiar future, unlike anything that has ever happened in the world before. It is one in which we organisms are all interacting like mad, in reshaping all of our environment. In the last few years we have stepped up by many orders of magnitude: in terms of communication speed, by a factor of 10^7 over what it was 100 years ago; in terms of travel speed, by a factor of 10^2 over 100 years ago; in weapons, by a factor of 10^6 over even 30 years ago, and so on. We have stepped into a new world of resources, of power, and of human interaction, within which evolution will take place on entirely different principles, with entirely different parameters of pressure and change from what has ever been before. As Lederberg has said, "There will be more evolution in the next 50 years than in the last 100,000."

This is the end of the era of evolution by natural selection. It is the beginning of evolution by human selection. Every species on the globe is going to be increasingly perturbed in its numbers by our predation, our pollution, our breeding, or our protection. The numbers will not be determined by what we used to call "natural selection" so much as by conscious or unconscious human intervention which is now shaping what the structure of the biosphere is going to be. In such an unpredictable large-scale crisis, we will be able to survive only by making maximum use of all we know about anticipatory problem-solving, with anticipatory control and anticipatory feedbacks. The danger is clear; only in this way can we keep from destroying our biosphere and ourselves, by collective accident. If we do survive, these feedbacks, of course, will necessarily lead to higher degrees of cooperation and lower degrees of conflict in the world. Our increasing interaction-intensity will necessarily lead to a new kind of mutuality in our decisions, a mutuality in which we can no longer treat each other as objects, but must treat each other as co-participants and co-subjects in all our experiments and plans. What this all means is that in the times just ahead, if we survive, we are going to move toward something much more like an aggregate decision-system—more like a single organism—then we have ever been before.

Perhaps it was implicit in the biological material and the physical necessities all along—as surely as the eventual development of wings is implicit in an atmosphere.