Access provided by School of the Art Institute of Chicago

| Muse | Search | Journals | This Journal | Contents |

Boundary Disputes: Homeostasis, Reflexivity, and the Foundations of Cybernetics

N. Katherine Hayles

Virtual reality did not spring, like Athena from the forehead of Zeus, full-blown from the mind of William Gibson. It has encoded within it a complex history of technological innovations, conceptual developments, and metaphorical linkages that are crucially important in determining how it will develop and what it is taken to signify. This essay explores that history by focusing on certain developments within cybernetics from the immediate post-World War II period to the present. These developments can be understood as progressing in three waves. The first period, 1945-1960, marks the foundational stage during which cybernetics was forged as an interdisciplinary framework that would allow humans, animals, and machines to be constituted through the common denominators of feedback loops, signal transmission, and goal-seeking behavior. The forum for these developments was a series of conferences sponsored by the Josiah Macy Foundation between 1946 and 1953. [1] Through the Macy discussions and the research presented there, the discipline solidified around key concepts and was disseminated into American intellectual communities by Macy **[End Page 441]** conferees, guests, and fellow travelers. Humans and machines had been equated for a long time, but it was largely through the Macy conferences that both were understood as information-processing systems.

Although space will not permit me to develop the second and third waves in detail, a synopsis will be useful in understanding the connections between them and later developments in virtual reality. The second wave was initiated by Heinz von Foerster, an Austrian émigré who became the coeditor of the Macy transcripts. This phase can be dated from 1960, when a collection of von Foerster's essays appeared under the title *Observing Systems*. [2] "Second-order cybernetics," von Foerster called the models he presented in these essays, because they extended cybernetic principles to the cyberneticians who had devised the principles. As his punning title recognizes, the observer of systems can himself be constituted as a system to be observed. The second wave reached its mature phase with Humberto Maturana's *Autopoiesis and Cognition*, coauthored with Francisco Varela. [3] Maturana and Varela developed von

Foerster's self-reflexive emphasis into a radical epistemology that saw the world as a set of formally closed systems. Organisms respond to their environment in ways determined by their internal self-organization. Hence they are not only self-organizing, they are also autopoietic, or self-making. Through the work of Maturana, Varela, and such theorists as the German sociologist Niklas Luhmann, [4] by 1980 cybernetics had moved from an umbrella term covering a collection of related concepts to a coherent theory undergirding the claims of a sophisticated and controversial epistemology.

The third wave emerged when virtual reality, combining the idea of a feedback loop with the greatly increased power of microcomputers, began to enter professional and consumer marketplaces with immersive devices that spliced the user's sensorium into three-dimensional simulated worlds. Instantiating the formal closure that Maturana saw as the basis for all perception in the closed loop that runs from the user's body through the simulation, VR technologies **[End Page 442]** provided a phenomenology to go along with the epistemology of autopoiesis. In the third wave, the idea of a virtual world of information that coexists with and interpenetrates the material world of objects is no longer an abstraction. It has come home to roost, so to speak, in the human sensorium itself.

Half a century beyond the watershed of the Macy conferences, feedback loops have become household words and cybernouns are breeding like flies, spawning cybernauts, cyberfutures, and cybersluts. People no longer find it strange to think of material objects as informational patterns. Hans Moravec, head of the Carnegie-Mellon Mobile Robot Laboratory, has even proposed a visionary scenario in which human consciousness is transformed into an informational pattern, extracted from the brain, and downloaded into a computer. [5] If there was ever a case to be made for a paradigm shift, this would seem to be it. Yet a close examination of threads in this tangled skein does not bear out either Kuhn's model of a revolutionary break or Foucault's vision of a new episteme springing suddenly into being across a wide range of cultural sites. Rather, the changes appear to be much more like what archeological anthropologists find when they study material culture: instead of a sudden break, change came about through overlapping patterns of innovation and replication. Part of my argument is thus historical, focusing on the specificities of how change occurred during the foundational period of cybernetics. Part of it is synthetic and conceptual, concerned to develop a schematic that will show how the three waves relate to each other (see Figure 1). The schematic illustrates the traffic between ideas and artifacts, suggesting that concepts, like artifacts, embody sedimented histories that affect how they develop. To make this argument, I will introduce a triad of terms appropriated from studies of material culture: seriation, skeuomorphs, and conceptual constellations. Finally, I will use this model of historical change to interrogate why it seems possible, at this cultural moment, for virtuality to displace materiality. I will argue that the claim for the self-sufficiency of a virtual world of information is deeply problematic, for this illusion, like everything that exists, has its basis in the very materiality it would deny.

Seriation, Skeuomorphs, and Conceptual Constellations

In archeology, changes in artifacts are customarily mapped through seriation charts. One may construct a seriation chart by parsing an artifact as a set of attributes that change over time. One **[End Page 443] [Begin Page 445]** of the attributes of a lamp, for example, is the element that gives off light. The first lamps used wicks for this

element. Later, with the discovery of electricity, wicks gave way to filaments. The figures that customarily emerge from this kind of analysis are shaped like a tiger's iris-- narrow at the top when an attribute first begins to be introduced, with a bulge in the middle during an attribute's heyday, tapering off at the bottom as the shift to a new model is completed. On a seriation chart for lamps, a line drawn at 1890 would show the figure for wicks waxing large, while the figure showing filaments would be intersected at the narrow tip of the top end; fifty years later, the wick figure would be tapering off, while the filament figure would be widening into its middle section. Considered as a set, the figures depicting changes in an artifact's attributes reveal patterns of overlapping innovation and replication. Some attributes are changed from one model to the next, while others remain essentially the same.

As Figure 1 suggests, the conceptual shifts that took place during the development of cybernetics display a pattern of overlapping innovation and replication reminiscent of material changes in artifacts. It is not surprising that conceptual fields should evolve similarly to material culture, for concept and artifact engage each other in continuous feedback loops. An artifact materially expresses the concept it embodies, but the process of its construction is far from passive. A glitch has to be fixed, a material exhibits unexpected properties, an emergent behavior surfaces--any of these challenges can give rise to a new concept, which results in another generation of artifact, which leads to the development of still other concepts. The reasoning suggests that it should be possible to construct a seriation chart tracing the development of a conceptual field analogous to seriation charts for artifacts.

In the course of the Macy conferences, certain ideas came to be associated with each other. Through a cumulative process that continued across several years of discussions, they were seen as mutually entailing each other until, like love and marriage, it seemed natural to the participants that they should go together. Such a constellation is the conceptual entity corresponding to an artifact, possessing an internal coherence that defines it as an operational unit. Its formation marks the beginning of a period; its disassembly and reconstruction, the transition to a different period. Indeed, periods are recognizable as such largely because constellations possess this coherence. Rarely is a constellation discarded wholesale; rather, some of the ideas comprising it are discarded, others are modified, and new ones are introduced. Like the attributes that comprise an **[End Page 445]** artifact, the ideas in a constellation change in a patchwork pattern of old and new.

During the Macy conferences, two constellations formed that were in competition with one another. One of these was deeply conservative, privileging constancy over change, predictability over complexity, equilibrium over evolution. At the center of this constellation was the concept of homeostasis, defined as the ability of an organism to maintain itself in a stable state. The other constellation led away from the closed circle of corrective feedback, privileging change over constancy, evolution over equilibrium, complexity over predictability. The central concept embedded in it was reflexivity, which for our purposes can be defined as turning a system's rules back on itself so as to cause it to engage in more complex behavior. In broader social terms, homeostasis reflected the desire for a "return to normalcy" after the maelstrom of World War II. By contrast, reflexivity pointed toward the open horizon of an unpredictable and increasingly complex postmodern world.

Around these two central concepts accreted a number of related ideas, like clumps of barnacles around the first mollusks brave enough to fasten themselves to a rock. Because the two constellations were in competition, each tended to define itself as what the other was not. They were engaged, in other words, in what Derrida has called an economy of supplementarity: the necessity *not* to be what the other *was* helped to define each partner in the dialectic. The reflexivity constellation, more amorphous and fuzzily defined than was homeostasis during the Macy period, finally collapsed as a viable model. At that point its homeostatic partner could not maintain itself in isolation, and out of this chaos a new constellation began to form through a seriated pattern of overlapping innovation and replication, taking elements from both of its predecessors and adding new features as well. This diachronic movement out of the homeostasis/reflexivity dialectic marks the end of the first wave of cybernetics.

Here I want to introduce another term from archeological anthropology. A skeuomorph is a design feature, no longer functional in itself, that refers back to an avatar that was functional at an earlier time. The dashboard of my Toyota Camry, for example, is covered by vinyl molded to simulate stitching; the simulated stitching alludes back to a fabric that was in fact stitched, although it no longer serves that function in my car. Skeuomorphs visibly testify to the social or psychological necessity for innovation to be tempered by replication. Like anachronisms, their pejorative first cousins, skeuomorphs are not unusual. On the contrary, they are so **[End Page 446]** deeply characteristic of the way concepts and artifacts evolve that it takes a great deal of conscious effort to avoid them. At Siggraph 93, the annual conference on computer graphics that showcases new products, I saw more skeuomorphs than morphs.

Perhaps the wittiest of these skeuomorphs was the "Catholic Turing Test" simulation, complete with a bench where the supplicant could kneel while making a confession by choosing selections from the video screen. [6] How can I understand the pleasure I took in this display? On one level, the installation alluded to the triumph of science over religion, for the role of divinely authorized interrogation and absolution had been taken over by a machine algorithm. On another level, the installation pointed to the intransigence of conditioned behavior, for the machine's form and function were determined by its religious precedessor. Like a Janus figure, the skeuomorph can look to both past and future, simultaneously reinforcing and undermining both. It calls into a play a psychodynamic that finds the new more acceptable when it recalls the old that it is in the process of displacing, and the traditional more comfortable when it is presented in a context that reminds us we can escape from it into the new. In the history of cybernetics, skeuomorphs acted as threshold devices, smoothing the transition between one conceptual constellation and another. Homeostasis, a foundational concept during the first wave, functions during the second wave as a skeuomorph. Although it is carried over into the new constellation, it ceases to be an initiating premise and instead performs the work of a gesture or an allusion used to authenticate the new elements in the emerging constellation. At the same time, it also exerts an inertial pull on the new elements that limits how radically they can transform the constellation.

A similar phenomenon appears in the transition from the second to the third wave. Reflexivity, the key concept of the second wave, is displaced in the third wave by emergence. Like homeostasis, reflexivity does not altogether disappear but lingers on as an allusion that authenticates new elements. It performs a more complex role than

mere nostalgia, however, for it also leaves its imprint on the new constellation in the possibility that the virtual space can close on itself to create an autopoietic world sufficient unto itself, independent of the larger reality in which it is embedded. The complex story formed by these seriated changes between the three **[End Page 447]** waves begins when humans and machines are equated by defining both as information-processing systems. If information is the link connecting humans and machines, then how it is defined is crucial. It is through this struggle that the first set of constellations is forged.

The Meaning of Information

From the beginning, "information" was a contested term. At issue was whether it should be defined solely as a mathematical function, without reference to its meaning to a receiver, or whether it should be linked to the context in which it is received and understood. The principal spokesperson for the mathematical definition was Claude Shannon, the brilliant electrical engineer at Bell Laboratories who formulated information as a function of the probability distribution of the elements comprising the message. [7] Shannon went out of his way to point out that this definition of information divorces it from the common understanding of information as conveying meaning. Information theory as Shannon developed it has only to do with the efficient transmission of messages through communication channels, not with what messages mean. Although others were quick to impute larger linguistic and social implications to his theory, Shannon was reluctant to generalize beyond the specialized denotation that information had within the context of electrical engineering. Responding to a presentation by Alex Bavelas on group communication at the eighth Macy conference, Shannon cautioned that he did not see "too close a connection between the notion of information as we use it in communication engineering and what you are doing here . . . the problem here is not so much finding the best encoding of symbols . . . but, rather, the determination of the semantic question of what to send and to whom to send it." [8] For Shannon, defining information was a strategic choice that enabled him to bracket semantics. He especially did not want to get involved in having to consider the receiver's mindset as part of the communication system. So strongly did he feel on the point that he suggested Bavelas distinguish between information in a channel and information in a human mind by characterizing the latter through "subjective probabilities," although how these were to be defined and calculated was by no means clear.

Donald MacKay took the opposite view, using his privileges as a **[End Page 448]** guest lecturer at the Macy conferences to argue for a close connection between information and meaning. In the rhetoric of the conferences, "objective" was associated with being "scientific," whereas "subjective" was a code word implying that one had fallen into a morass of unquantifiable feelings that might be magnificent but were certainly not science. MacKay's first move was to rescue information affecting the receiver's mindset from the "subjective" label. He proposed that both Shannon and Bavelas were concerned with what he called "selective information"--that is, information calculated by considering the selection of message elements from a set. MacKay argued for another kind of information that he called "structural": structural information has the capacity to "increase the number of dimensions in the information space" by acting as a metacommunication. [9]

To illustrate the distinction, say I launch into a joke and it falls flat; in that case, I may resort to telling my interlocutor, "That was a joke." The joke's message content, considered as selective information, can be calculated as a function of the probabilities of the message elements. Performing this calculation is equivalent to operating within a two-dimensional space, for the only variables are the probability and probability function. By contrast, my comment identifying the message as a joke is structural information, for it implies the existence of two different kinds of messages--jokes and serious statements. To accommodate this distinction, another parameter is necessary. This parametric variation can be represented by stacking two planes on top of one another, with the vertical dimension expressing the relation between the two types of information. Other kinds of metacommunications--for example, comments distinguishing between literal and metaphoric statements--would add additional parameters and more stacks. In another image that MacKay liked to use, he envisioned selective information as choosing among folders in a file drawer, whereas structural information increased the number of drawers.

Since structural information amounts to information on how to interpret a message, the effect on the receiver necessarily enters the picture. By calling such information "structural" rather than "subjective," MacKay changed its connotation. Not only did "structural" remove the suggestion of unscientific subjectivity, it also elevated the noun it modified to the status of a metacommunication that controls the subsystem with which it communicates. Thus the **[End Page 449]** move from "subjective" to "structural" information made a negatively encoded term into a positive position of power. In a sense, of course, he was doing no more than any good rhetorician would-- choosing his words carefully so they produced the desired effect in his audience. But entwined with this rhetorical effect was a model that triangulated between reflexivity, information, and meaning. Arguing for a strong correlation between the *nature* of a representation and its *effect*, his argument recognized the mutual constitution of language and content, message and receiver.

The problem was how to quantify the model. It implied that representations have a double valence: seen in one perspective, they are measuring instruments that point out into the world; in another perspective, they point back to the human agents who created them. By registering changes in the measuring instruments, one can infer something about the mental states of the agents who made them. And how does one perform this calculation? Through changes in an observing mind, which in turn can also be observed and measured by another mind. The progression inevitably turns into the infinite regress characteristic of reflexivity. This kind of move is familiar in twentieth-century art and literature, from the drawings of M. C. Escher to the fictions of Borges. Indeed, it is hardly an exaggeration to say that it has become an artistic and critical commonplace. Finding it in a quantitative field like information theory is more unusual. In the context of the Macy conferences, MacKay's conclusion qualified as radical: reflexivity, far from being a morass to be avoided, is precisely what enables information and meaning to be connected.

To achieve quantification, however, it was necessary to have a mathematical model for the changes the message triggered in the receiver's mind. The staggering problems this presented no doubt explain why MacKay's version of information theory was not widely accepted among the electrical engineers who would be writing, reading, and teaching the textbooks on information theory in the coming decades. Historically it was

Shannon's definition of information, not MacKay's, that became the industry standard. The issues underlying the debate between Shannon and MacKay are important for the Macy conferences, for they instantiate a central problem faced by the participants. Throughout the transcripts, there is constant tension between the desire to define problems narrowly so that reliable quantification could be achieved, and the desire to define them broadly so that they could address the momentous questions that kept edging their way into the discussions. These conflicting desires kept getting tangled up with each other. On the one hand, broad implications were drawn from narrowly **[End Page 450]** constructed problems. On the other, problems constructed so broadly that they were unworkable quantitatively were nevertheless treated as if they were viable scientific models. The discrepancy led Steve Heims, in his study of the Macy conferences, to remark that much of the "so-called social science was unconvincing to me as science in any traditional sense. In fact, some of it seemed to have only a thin scientific veneer, which apparently sufficed to make it acceptable." [10]

My interest in this tension has a different focus. I want to show how it works as a metaphoric exchanger to construct the human in terms of the mechanical, and the mechanical in terms of the human. Precisely because there was continuing tension between quantification and implication, passages were established between human intelligence and machine behavior. As the rival constellations of homeostasis and reflexivity began to take shape, man and machine were defined in similar terms--as homeostats, or as reflexive devices that threatened to draw the scientists into a morass of subjectivity. Three tropes will illustrate how these processes of metaphoric exchange were mediated by the traffic between concept and artifact. The first crossroads for this material/conceptual traffic is Shannon's electronic rat, a goal-seeking machine that modeled a rat learning a maze. The second is Ross Ashby's homeostat, a device that sought to return to an equilibrium state when disturbed. The third, more an image than a model, envisions a man spliced into an information circuit between two machines. By most standards, for example those invoked by Steve Heims when he questions whether the Macy presentations were good science, the electronic rat and the homeostat were legitimate scientific models. By contrast, the man-in-the-middle image was so loosely defined that it can scarcely qualify as a model at all; moreover, by the time of the Macy conferences it had become a skeuomorph, for automatic devices that replaced the man-in-the-middle had already been used in World War II as early as 1942. [11] Nevertheless, all three tropes functioned **[End Page 451]** as exchangers that brought man and machine into equivalence; all shaped the kinds of stories that participants would tell about what this equivalence meant. In this sense it is irrelevant that some were "good science" and some were not, for they were mutually interactive in establishing the presuppositions of the field. As much as any of the formal theories, these presuppositions defined the shape of the first wave of cybernetics and influenced its direction of flow.

The Electronic Rat, the Homeostat and the Man-in-the-Middle

There are moments of clarity when the participants in the Macy conferences came close to articulating explicitly the presuppositions informing the deep structure of the discussion. At the seventh conference, John Stroud of the U.S. Naval Electronic Laboratory in San Diego pointed to the far-reaching implications of Shannon's construction of information through the binary distinction between signal and noise.

"Mr. Shannon is perfectly justified in being as arbitrary as he wishes," Stroud observed:

> We who listen to him must always keep in mind that he has done so. Nothing that comes out of rigorous argument will be uncontaminated by the particular set of decisions that were made by him at the beginning, and it is rather dangerous at times to generalize. If we at any time relax our awareness of the way in which we originally defined the signal, we thereby automatically call all of the remainder of the received message the "not" signal or noise. [12]

As Stroud realized, Shannon's distinction between signal and noise had a conservative bias that privileges stasis over change. Noise interferes with the exact replication of the message, which is presumed to be the desired result. The structure of the theory implied that change was deviation, and that deviation should be corrected. By contrast, MacKay's theory had as its generative distinction the difference in the state of the receiver's mind before and after the message arrived. In his model, information was not *opposed* to change, it *was* change.

Applied to goal-seeking behavior, the two theories pointed in different directions. Privileging signal over noise, Shannon's theory implied that the goal was a preexisting state toward which the mechanism would move by making a series of distinctions between correct and incorrect choices. The goal was stable, and the mechanism would achieve stability when it reached the goal. This construction **[End Page 452]** easily led to the implication that the goal, formulated in general and abstract terms, was less a specific site than stability itself. Thus the construction of information as a signal/noise distinction and the privileging of homeostasis produced and were produced by each other. By contrast, MacKay's theory implied that the goal was not a fixed point but a constantly evolving dance between expectation and surprise. In his model, setting a goal temporarily marked a state that itself would become enfolded into a reflexive spiral of change. In Gregory Bateson's phrase, information was a difference that made a difference. In the same way that signal/ noise and homeostasis went together, so did reflexivity and information as a signifying difference.

These correlations imply that before Shannon's electronic rat ever set marker in maze, it was constituted through assumptions that affected how it would be interpreted. Although Shannon called his device a maze-solving machine, the Macy group quickly dubbed it a rat. [13] The machine consisted of a 55 square grid, through which a sensing finger moved. An electric jack that could be plugged into any of the 25 squares marked the goal, and the machine's task was to move through the squares by orderly search procedures until it reached the jack. The machine could remember previous search patterns and either repeat them or not, depending on whether they had been successful. While Heinz von Foerster, Margaret Mead, and Hans Teuber in their introduction to the eighth conference volume highlighted the electronic rat's significance, they also acknowledged its limitations: "We all know that we ought to study the organism, and not the computers, if we wish to understand the organism. Differences in levels of organization may be more than quantitative." [14] They went on to argue, however, that "the computing robot provides us with analogues that are helpful as far as they seem to hold, and no less helpful whenever they break down. To find out in what ways a nervous system (or a social group) differs from our man-made analogues requires experiment. These experiments would not have been considered if

the analogue had not been proposed." [15]

There is another way to understand this linkage. By suggesting certain kinds of experiments, the analogues between intelligent machines and humans *construct the human in terms of the machine.* **[End Page 453]** Even when the experiment fails, the basic terms of the comparison operate to constitute the signifying difference. If I say a chicken is not like a tractor, I have characterized the chicken in terms of the tractor, no less than when I assert that the two are alike. In the same way, whether it is understood as like or unlike, human intelligence ranged alongside an intelligent machine is put into a relay system that constitutes the human as a special kind of information machine, and the information machine as a special kind of human. Moreover, although some characteristics of the analogy may be explicitly denied, the presuppositions it embodies cannot be denied, for they are intrinsic to being able to think the model. The presuppositions embodied in the electronic rat include the idea that both humans and cybernetic machines are goal-seeking mechanisms learning through corrective feedback to reach a stable state. Both are information processors that tend toward homeostasis when they are functioning correctly.

Given these assumptions, it is perhaps predictable that reflexivity should be constructed in this model as neurosis. Shannon, demonstrating how his electronic rat could get caught in a reflexive loop that would keep it circling endlessly around, remarked that "it has established a vicious circle, or a singing condition." [16] "Singing condition" is a phrase that Warren McCulloch and Warren Pitts had used in an earlier presentation to describe neuroses modeled through cybernetic neural nets. If machines are like humans in having neuroses, humans are like machines in having neuroses that can be modeled mechanically. Linking humans and machines in a common circuit, the analogy constructs both as equilibrium systems that become pathological when they fall into reflexivity. This kind of mutually constitutive interaction belies the implication in the volume's introduction that such analogues are neutral heuristic devices. More accurately, they are relay systems that transport assumptions from one arena to the next. Some of these assumptions may be explicitly recognized and in this sense authorized; others are not. Whether authorized or unauthorized, they are part of the context that guides inquiry, suggests models, and intimates conclusions.

The assumptions traveling across the relay system set up by homeostasis are perhaps most visible in the discussion of W. Ross Ashby's homeostat. [17] The homeostat was an electrical device constructed **[End Page 454]** with transducers and variable resistors. When it received an input changing its state, it searched for the configuration of variables that would return it to its initial condition. Ashby explained that the homeostat was meant to model an organism that must keep essential variables within preset limits in order to survive. He emphasized that the cost of exceeding those limits is death: if homeostasis equals safety ("Your life would be safe," Ashby responded, when demonstrating how the machine could return to homeostasis [18]), departure from homeostasis threatens death. It is not difficult to discern in this rhetoric of danger and safety echoes of the traumatic experiences of World War II. One of Ashby's examples, for instance, concerns an engineer sitting at the control panel of a ship: the engineer functions like a homeostat as he strives to keep the dials within certain limits to prevent catastrophe. Human and machine are alike in needing stable interior environments. The human keeps the ship's interior stable, and this stability preserves the

homeostasis of the human's interior, which in its turn allows the human to continue to ensure the ship's homeostasis. Arguing that homeostasis is a requirement "uniform among the inanimate and the animate," Ashby privileged it as a universally desirable state. [19]

The postwar context for the Macy conferences played an important role in formulating what counted as homeostasis. Given the cataclysm of the war, it seemed self-evident that homeostasis was meaningful only if it included the environment as part of the picture. Thus Ashby conceived of the homeostat as a device that included both the organism and the environment. "Our question is how the organism is going to struggle with its environment," he remarked, "and if that question is to be treated adequately, we must assume some specific environment." [20] This specificity was expressed through the homeostat's four units, which could be arranged in various configurations to simulate organism-plus-environment. For example, one unit could be designated "organism" and the remaining three the "environment"; in another arrangement, three of the units might be the "organism," with the remaining one the "environment." Formulated in general terms, the problem the homeostat addressed was this: Given some function of the environment E, can the organism find an inverse function E-1 such that the product of the two will result in an equilibrium state? **[End Page 455]** When Ashby asked Macy participants whether such a solution could be found for highly nonlinear systems, Julian H. Bigelow correctly answered "In general, no." [21] Yet, as Walter Pitts observed, the fact that an organism continues to live means that a solution does exist. More precisely, the problem was whether a solution could be articulated within the mathematical conventions and technologies of representation available to express it. These limits, in turn, were constituted through the specificities of the model that translated between the question in the abstract and the particular question posed by that experiment. Thus the emphasis shifted from finding a solution to stating the problem.

This dynamic appears repeatedly through the Macy discussions. Participants increasingly understood the ability to specify exactly what was wanted as the limiting factor for building machines that could perform human functions. At the ninth conference, Walter Pitts was confident enough of the construction to claim it as accepted knowledge: "At the very beginning of these meetings, the question was frequently under discussion of whether a machine could be built which would do a particular thing, and, of course, the answer, which everybody has realized by now, is that as long as you definitely specify what you want the machine to do, you can, in principle, build a machine to do it." [22] If what is exactly stated can be done by a machine, the residue of the uniquely human becomes coextensive with the qualities of language that interfere with precise specification--its ambiguity, metaphoric play, multiple encoding, and allusive exchanges between one symbol system and another. The uniqueness of human behavior thus becomes assimilated to the ineffability of language, while the common ground that humans and machines share is identified with the univocality of an instrumental language that has banished ambiguity from its lexicon. This train of thought indicates how the rival constellations of homeostasis and reflexivity assimilated other elements into themselves. On the side of homeostasis was instrumental language, while ambiguity, allusion, and metaphor stood with reflexivity.

By today's standards Ashby's homeostat was a simple machine, but it had encoded

within it a complex network of assumptions. Paradoxically, the model's simplicity facilitated rather than hampered the overlay of assumptions onto the artifact, for its very lack of complicating detail meant that the model stood for much more than it physically enacted. Ashby acknowledged during discussion **[End Page 456]** that the homeostat was a simple model and asserted that he "would like to get on to the more difficult case of the clever animal that has a lot of nervous system and is, nevertheless, trying to get itself stable." [23] The slippage between the simplicity of the model and the complexity of the phenomena did not go unremarked. J. Z. Young, from the anatomy department at University College London, sharply responded, "Actually that is experimentally rather dangerous. You are all talking about the cortex and you have it very much in mind. Simpler systems have only a limited number of possibilities." [24] Yet the "simpler systems" helped to reinforce the idea that humans are mechanisms that respond to their environments by trying to maintain homeostasis; that the function of scientific language is exact specification; that the bottleneck for creating intelligent machines lay in formulating problems exactly; and that a concept of information that privileges exactness over meaning is therefore more suitable to model construction than one that does not. Ashby's homeostat, Shannon's information theory, and the electronic rat were collaborators in constructing an interconnected network of assumptions about language, teleology, and human behavior.

These assumptions did not go uncontested. The concept that most clearly brought them into question was reflexivity. Appropriately, the borderland where reflexivity contested the claims of homeostasis was the man-in-the-middle. The image was introduced in the sixth conference in John Stroud's analysis of an operator sandwiched between a radar tracking device on one side and an antiaircraft gun on the other. The gun operator, Stroud observed, is "surrounded on both sides by very precisely known mechanisms and the question comes up, 'What kind of a machine have we put in the middle.'" [25] The image as Stroud used it constructs the man as an input/output device: information comes in from the radar, travels through the man, and goes out through the gun. The man is significantly placed in the *middle* of the circuit, where his output and input are already spliced into an existing loop. Were he at the end, it might be necessary to consider more complex factors, such as how he was interacting with an open-ended evironment whose state could not be exactly specified and whose future evolution was unknown. The focus in Stroud's presentation was on how information **[End Page 457]** is transformed as it moves through the man-in-the-middle. As with the electronic rat and the homeostat, the emphasis was on predictability and homeostatic stability.

Countering this view was Frank Fremont-Smith's recognition of the inescapable reflexivity inherent in constructing this system as a system. "Probably man is never only between the two machines," Fremont-Smith pointed out. "Certainly he is never only in between two machines when you are studying him because you are the other man who is making an input into the man. You are studying and changing his relation to the machines by virtue of the fact that you are studying him." [26] Fremont-Smith's opening of the circuit to the environment through reflexivity was later countered by Stroud in a revealing image that sought once again to close the circuit: "The human being is the most marvelous set of instruments," he observed, "but like all portable instrument sets the human observer is noisy and erratic in operation. However, if these are all the instruments you have, you have to work with them until something better

comes along." [27] In Stroud's remark, the open-endedness of Fremont-Smith's construction is converted into a portable instrument set. The instrument may not be physically connected to two mechanistic terminals, the image implied, but this lack of tight connection only makes the splice invisible. It does not negate the suture that constructs the human as an information-processing machine that ideally should be homeostatic in its operation, however noisy it is in practice.

As his switch to formal address indicates, Fremont-Smith was apparently upset at the recuperation of his comment back into the presuppositions of homeostasis. "You cannot possibly, Dr. Stroud, eliminate the human being. Therefore what I am saying and trying to emphasize is that, with all their limitations, it might be pertinent for those scientific investigators at the general level, who find to their horror that we have to work with human beings, to make as much use as possible of the insights available as to what human beings are like and how they operate." [28] His comment cuts to the heart of the objection against reflexivity. Whether construed as "subjective information" or as changes in the observer's representations, reflexivity opens the man-in-the-middle to internal psychological complexity so that he can no longer be constructed as a black box functioning as an input/output device. The fear is that **[End Page 458]** under these conditions, reliable quantification becomes elusive or impossible and science slips into subjectivity, which to many conferees meant that it was not real science at all. Confirming traditional ideas of how science should be done in a postwar atmosphere that was already clouded by the hysteria of McCarthyism, homeostasis thus implied a return to normalcy in more than one sense.

The thrust of Fremont-Smith's observations was, of course, to intimate that psychological complexity was unavoidable. The responses of other participants reveal that it was precisely this implication they were most concerned to deny. Were their responses valid objections, or themselves evidence of the very subconscious resistance they were trying to disavow? The primary spoksperson for this latter disconcerting possibility was Lawrence Kubie, a psychoanalyst from the Yale University Psychiatric Clinic. In correspondence, Kubie enraged other participants by interpreting their comments as evidence of their psychological states rather than as matters for scientific debate. In his presentations he was more tactful, but the reflexive thrust of his argument remained clear. His presentations occupy more space in the published transcripts than those of any other participant, comprising about one-sixth of the total. Although he met with repeated skepticism among the physical scientists, he continued to try to explain and defend his position. At the center of his explanation was the multiply encoded nature of language, operating at once as an instrument that the speaker could use to communicate and as a reflexive mirror that revealed more than the speaker knew. Like MacKay's theory of information, Kubie's psychoanalytic approach built reflexivity into the model. Also like MacKay's theory, the greatest conscious resistance it met was the demand for reliable quantification.

Kubie's presentations grew increasingly entrenched as the conferences proceeded. The resistance they generated illustrates why reflexivity had to be redefined if it was to be rescued from the dead end to which, in the view of many participants, it seemed to lead. From the point of view of those who resisted Kubie's ideas, psychoanalysis collapsed the distance between speaker and language, turning what should be objective scientific debate into a tar baby that clung to them more closely the more they tried to push it away. The association of reflexivity with psychoanalysis probably

delivered the death blow to the reflexivity constellation. Homeostasis seemed to have won the day. Ironically, however, its triumph was also its senescence, for, divorced from the reflexivity constellation, it lost its power to generate new ideas and to serve as an impetus **[End Page 459]** for further research. After about 1960, homeostasis became a skeuomorph, pointing back to an earlier period but also serving as a link to a more radical form of reflexivity. To be acceptable to the community that grew out of the Macy conferences, a mode of reflexivity had to be devised that could be contained within stable boundaries and divorced from the feedback loop that implicated the observer in everything he said. Maturana's epistemology of autopoiesis satisfied these conditions in one sense, although, as we shall see, in another sense it profoundly subverted them.

Autopoiesis and the Closure of the System

Maturana's epistemology is grounded in work he did in the late 1950s and early 1960s on visual processing in the frog's cortex. He coauthored the influential paper "What the Frog's Eye Tells the Frog's Brain," which demonstrated that the frog's visual cortex responds to stimuli in ways specific to the species. [29] The frog's brain discussed in this article, far from being a "natural" object, is as much an artifact as the electronic rat and the homeostat. To make the brain productive of scientific knowledge, precisely placed microelectrodes were inserted into it using stereotactic surgical equipment; the electrodes were then connected to complex electronic monitoring systems. Spliced into a circuit that included the experimenter's observations as well as the laboratory equipment, the frog brain (I drop the possessive because at this point the brain no longer belonged strictly to the frog) became a techno-bioapparatus instantiating the cybernetic framework that constituted animals and machines as information-processing systems, even as it was used to advance and develop that same framework.

Thus reflexively constructed, the frog brain showed little response to large slow-moving objects, while small quickly-moving objects evoked a large response. The response pattern is obviously adaptive for frogs, since it enables them to detect flies and other prey while screening out other phenomena that might distract them. The radical implication of this work is summarized in the authors' conclusion that "the [frog's] eye speaks to the brain in a language already highly processed and interpreted." [30] Imaginatively giving the frog brain back to the frog, I have elsewhere tried to envision what Newton's three laws of motion might look like to a **[End Page 460]** frog. The commonsense observation that an object at rest remains the same object when it is in motion would be scarcely conceivable from a frog's point of view, for the two kinds of objects are processed in completely different ways by the frog's sensory system. The point is not, of course, that humans can see more or better than frogs. Rather, if perception is species-specific, then it follows that every perception is always already encoded by the perceptual apparatus of the observer, whether the observer is a human or a frog. Thus there is no possibility of a transcendent position from which to see reality as it "really" is. Simply put, the article blows a frog-sized hole in objectivist epistemology.

Despite its radical implications, the article's *form* is in the mainstream of scientific tradition, for nothing in the way it is written challenges scientific objectivity. Nowhere do the authors acknowledge that their observations are made from a human

perspective and thus are relative to the perceptual apparatus of the human sensorium. Later work by Maturana made this inconsistency even more apparent, for it questioned whether the [primate] brain creates a "representation" of the outside world at all. [31] Data from the techno-bioapparatus of the primate visual cortex demonstrated that there is no qualitative correspondence between stimulus and response, and only a small quantitative correspondence. Maturana thus concluded that the stimulus acts as a trigger for a response dictated almost entirely by the internal organization of the frog sensory receptors and central nervous system. But here language betrays me, for if what happens inside the frog cortex is not a representation of what happens outside, then it is misleading to talk about stimulus and response, for such language implies a one-to-one correlation between the two events. What is needed, evidently, is another kind of language that would do justice to Maturana's revolutionary insight that "there is no observation without an observer."

To solve the problem, Maturana proposed a radical new epistemology that rejected traditional causality. So powerful is this epistemology that it can rightfully be considered to create a different kind of world, which I will attribute to Maturana (thus ironically restoring to his name the possessive that had earlier been taken from the frog's brain when it became a site for scientific knowledge production). In Maturana's world, one event does not cause another **[End Page 461]** --rather, events act as "triggers" for responses determined by a system's self-organization. Maturana defined a self-organizing system as a composite unity: it is a unity because it has a coherent organization, and it is composite because it consists of components whose relations with each other and with other systems constitute the organization that defines the system as such. Thus the components constitute the system, and the system unites the components. The circularity of the reasoning foregrounds reflexivity while also transforming it. Whereas in the Macy conferences reflexivity was associated with psychological complexity, in Maturana's world it is constituted through the interplay between a system and its components. They mutually define each other in the bootstrap operation characteristic of reflexive self-constitution.

Reflexivity is also central to the distinction Maturana makes between allopoietic and autopoietic systems. Allopoietic systems have as their goal something exterior to themselves. When I use my car to drive to work, I am using it as an allopoietic system, for its function is transportation, a goal exterior to the maintenance of its internal organization. As the example indicates, allopoietic systems are defined functionally and teleologically rather than reflexively. By contrast, autopoietic systems have as their goal the maintenance of their own organization. If my foremost purpose in life is to continue living, then I am functioning as an autopoietic system, for I have as my goal the maintenance of my self-organization. One can see in this formulation the ghost of homeostasis, although it now signifies not so much stability as a formal closure of the system upon itself.

Maturana's world was developed in part as a reaction against behaviorism. Von Foerster, who worked with him on a number of projects, anticipated Maturana's rejection of behaviorism when he contested the behaviorist account of a conditioned subject as a "black box" that gives a predictable output for a known input (a scenario reminiscent of what Stroud wanted to do with the man-in-the-middle). Von Foerster turned behaviorism on its head by shifting the focus to the experimenter-observer. He argued that behaviorist experiments do not prove that living creatures are black boxes;

rather, they demonstrate that the experimenter has simplified his environment so it has become predictable, while perserving intact his own complexity and free will. [32] In Maturana's terms, the experimenter has converted the experimental subject into an allopoietic **[End Page 462]** system, while continuing to function himself as an autopoietic system. The critique gives a political edge to Maturana's epistemology, for it points to the power relations that determine who gets to function autopoietically and who is reduced to allopoiesis. Applying von Foerster's arguments to Maturana's experimental work leads to the ironic conclusion that the frog brain was made into an allopoietic system so that it could buttress arguments for the importance of autopoiesis.

The reasoning illustrates the difficulty of working out the implications of a reflexive epistemology, for the ground keeps shifting depending on which viewpoint is adopted (for example, that of the frog versus that of the experimenter). Maturana was able to carry his conclusions as far as he did because he displaced the focus of attention from the boundary between a system and the environment, to the feedback loops within the organism. The price he pays for the valuable insights this move yields is the erasure of the environment. In Maturana's world, the environment becomes a nebulous "medium" populated by self-organizing systems that interact with each other through their structural couplings with the medium. Resisting the closure into black boxes that the reductive causality of behaviorism would effect, Maturana performs another kind of closure that can give only a weak account of how systems interact with their environment and with each other. If reflexivity is finally given its due, it is at the price of giving a full and rich account of interactivity.

In the third wave of cybernetics, the self-organization that is a central feature of Maturana's world lingers on as a skeuomorph, and the emphasis shifts to emergence and immersion. Whereas for Maturana self-organization was associated with homeostasis, in the simulated computer worlds of the third wave, self-organization is seen as the engine driving systems toward emergence. Interest is focused not on how systems maintain their organization intact, but rather on how they evolve in unpredictable and often highly complex ways through emergent processes. Although I take these simulated worlds to include artificial life as well as virtual reality, in the interest of space I will discuss only the latter, and that briefly. How does the history of overlapping innovation and replication, of skeuomorphs connecting one era with another, and of traffic through such artifacts as the homeostat, the electronic rat, and the frog cortex, matter to the development of virtual reality? What webbed network of connections mediate cybernetics for virtual reality technologies, and what issues are foregrounded by exploring these connections? **[End Page 463]**

The Sedimented History of Virtual Reality

Cybernetics is connected to virtual reality technologies in much the same way as Cartesian space is connected to contemporary mapmaking. Through such seminal ideas as information, feedback loops, human-machine interfaces, and circular causality, cybernetics provided the terminology and conceptual framework that made virtual reality a possibility, although the pathways between any given cybernetic theory and virtual reality technology may be indirect and highly mediated. From this tangled web, I will pull three strands that I think have important consequences for VR: embodiment, reflexivity, and positionality. The first thread is spun from Claude Shannon's move of conceptualizing information as a pattern distinct from the physical

markers that embody it. As we have seen, this move allowed information to be reliably quantified. It also created a way of thinking about information that made it seem disembodied, removed from the material substrate in which it is instantiated. This construction of information allows cyberspace to be conceptualized as a disembodied realm of information that humans enter by leaving their bodies behind. In this realm, so the story goes, we are transformed into information ourselves and thus freed from the constraints of embodiment. We can take whatever form we wish, including no form at all.

In fact, of course, we are never disembodied. Simulated worlds can exist for us only because we can perceive them through the techno-bioapparatus of our body spliced into the cybernetic circuit. The reading of cyberspace as a disembodied realm is a skeuomorph that harks back to the first wave of cybernetics, which in turn is a reading of information that reinscribes into cybernetics a very old and traditional distinction between form and matter. These residues, echoing in a chain of allusion and reinscription that stretches back to Plato's cave, testify to the importance of excavating the sedimented history of artifacts and concepts, for they allow us to understand how the inertial weight of tradition continues to exert gravitational pull on the present.

Although the perception that cyberspace is disembodied is refuted by the material realities of the situation, it nevertheless has a material effect on what technologies will be developed, how they will be used, and what kind of virtual worlds they will instantiate. If the point is to enhance perceptions of disembodiment, then the technology will insulate the users as much as possible from their immediate surroundings. Rather than develop open systems such as Mandala that emphasize the user's connection to the environment, **[End Page 464]** the industry will continue to push head-mounted displays and stereo eyephones that cut the users off from their surroundings in the real world. If, on the contrary, the link between the virtual experience and embodiment is perceived to be important, then simulations such as *Placeholder* will be developed. [33] Designed by Brenda Laurel and Rachel Strickland, *Placeholder* requires the user to choose one of four totemic animals in which to be embodied. If the user chooses Crow, she negotiates the virtual terrain by flying; if Snake, her vision shifts to infrared and she moves by crawling. Obviously this is not a reinscription of the "natural" body in the virtual world, since these experiences are not normally available to the human sensorium. Rather, the simulation recognizes that the virtual body is a techno-bioapparatus, but in a way that emphasizes rather than conceals the centrality of embodiment to experience.

The next thread I want to pull from the skein is reflexivity. The struggle to introduce reflexivity into cybernetics makes clear how difficult it is to include the observer in the situation and to realize fully the implications of this inclusion. Virtual reality technologies can facilitate this realization, for by providing a prosthesis to the "natural" sensorium, they make the experience of mediated perception immediately obvious to the user. Structural coupling, Maturana's phrase for how self-organizing systems interact with each other and the surrounding medium, seems a cumbersome, roundabout way to say something simple like "I see the dog" if one is speaking from the position of the "natural" body. If the dog appears in a VR simulation, however, then it becomes common sense to realize that one "sees" the dog only through the structural couplings that put one's visual cortex in a feedback loop with the simulated image. These couplings include the interfaces between the retina and the stereovision

helmet, between the helmet and the computer through data transmission cables, between the incoming data and the CTR display via computer algorithms, and between the algorithms and silicon chips through the magnetic polarities recorded on the chips. In this sense VR technologies instantiate Maturana's world, converting what may seem like abstract and far-out ideas into experiential realities.

Taken to the extreme, the awareness of the mediated nature of **[End Page 465]** perception that VR technologies provide can be taken to signify that the body itself is a prosthesis. [34] The only difference between the body and VR couplings, in this view, is that the body was acquired before birth through organic means rather than purchased from some high-tech VR laboratory like VPL. The body, like the VR body-suit, creates mediated perceptions; both operate through structural couplings with the environment. In this account we can see the thread of disembodiment getting entangled with the thread of reflexivity, creating a view of the subject that sees human embodiment as one option among many, neither more nor less artificial than the VR prostheses that extend perception into simulated worlds. Is it necessary to insist, once again, that embodiment is not an option but a necessity in order for life to exist? The account elides the complex structures that have evolved through millennia in coadaptive processes that have fitted us as organisms to our environment. Compared to the complexity and depth of this adaptation, the VR body-suit is a wrinkle in the fabric of human life, hardly to be mistaken as an alternative form of embodiment. Knowing the history of cybernetics can be helpful in maintaining a sense of scale, for the move of mapping complex assumptions onto relatively simple artifacts has occurred before, usually in the service of eliding the immense differences between the complexities of the human organism and the comparatively simple architectures of the machines.

The last thread I want to unravel concerns the position of the VR user. I take "positionality" to include the body in which the user is incarnate, the language she speaks, and the culture in which she is immersed, as well as the specificities and collectivities of her individual history--but the part of positionality that concerns me here is her relation to the technology. The narrative of cybernetics as I have constructed it here suggests that the field is moving along a trajectory that arcs from homeostasis to reflexivity to emergence/immersion. First stability is privileged; then a system's ability to take as its goal the maintenance of its own organization; then its ability to manifest emergent and unpredictable properties. Inscribing the human subject into this trajectory, we can say that in the first stage, the privileged goal is for the human to remain an autonomous and homeostatic subject; in the second stage, to change structurally but nevertheless to maintain her internal organization **[End Page 466]** intact; and in the third stage, to mutate into a new kind of form through emergent processes that evolve spontaneously through feedback loops between human and machine.

The larger narrative inscribed here thus locates the subject in a changing relation to intelligent machines that points toward a looming transformation: the era of the human is about to give way, or has already given way, to the posthuman. There are already in circulation various accounts of how this transformation will come about and what it will mean. Howard Rheingold has called it IA, intelligence augmentation, arguing that humans and intelligent machines are entering into a symbiosis to which each will bring the talents and gifts specific to their species: humans will contribute to the partnership pattern recognition, language capability, and understanding ambiguities; machines will

contribute rapid calculation, massive memory storage, and rapid data retrieval. [35] Bruce Mazlish has called the posthuman era the fourth discontinuity, arguing that it constitutes the latest of four decisive breaks in human subjectivity. [36] Hans Moravec sees the break more pessimistically, arguing that protein-based life forms are about to be superseded by silicon-based life and that humans will soon become obsolete. [37] The differences between these accounts notwithstanding, they concur in seeing the posthuman era as constituting a decisive break in the history of humankind. The narrative offered here aims to counter this apocalyptic tone of sudden and irreversible change by looking closely at how change has actually occurred in the history of cybernetics. Concepts and artifacts are never invented out of whole cloth; rather, they embody a sedimented history that exerts an inertial pull on the new even as it modifies the old. If it is true that we are on the threshold of becoming posthuman, surely it behooves us to understand the overlapping patterns of replication and innovation that have brought us to where we now are.

University of California at Los Angeles

N. Katherine Hayles is Professor of English at the University of California at Los Angeles. She writes on literature and science in the twentieth century. Her books include *Chaos Bound: Orderly Disorder in Contemporary Literature and Science*, and *The Cosmic Web: Scientific Field Models and Literary Strategies in the Twentieth Century*. Currently she is completing a book entitled, *Virtual Bodies: Cybernetics, Literature, Information*.

Notes

1. Five of the Macy Conference transactions were published under the title *Cybernetics: Circular Causal and Feedback Mechanisms in Biological and Social Systems,* vols. 6-10, ed. Heinz von Foerster (New York: Macy Foundation, 1949-1955) (hereinafter cited as *Cybernetics*, with conference number and year). The best study of the Macy conferences is Steve J. Heims, *The Cybernetics Group* (Cambridge, Mass.: MIT Press, 1991); in addition to discussing the conferences, Heims also conducted interviews with many of the participants who have since died.

2. Heinz von Foerster, *Observing Systems,* 2nd ed. (Salinas: Intersystems Publications, 1984).

3. Humberto R. Maturana and Francisco J. Varela, *Autopoiesis and Cognition: The Realization of the Living*, Boston Studies in the Philosophy of Science, vol. 42 (Dordrecht: D. Reidel, 1980).

4. Niklas Luhmann has modified and extended Maturana's epistemology in significant ways; see, for example, his *Essays on Self-Reference* (New York: Columbia University Press, 1990) and "The Cognitive Program of Constructivism and a Reality that Remains Unknown," in *Selforganization: Portrait of a Scientific Revolution,* ed. Wolfgang Krohn et al. (Dordrecht: Kluwer Academic Publishers, 1990), pp. 64-85.

5. Hans Moravec, *Mind Children: The Future of Robot and Human Intelligence* (Cambridge, Mass.: Harvard University Press, 1988), pp. 109-110.

6. The simulation is the creation of Gregory P. Garvey of Concordia University. An account of it can be found in *Visual Proceedings: The Art and Interdisciplinary Programs of Siggraph 93,* ed. Thomas E. Linehan (New York: Association for Computing Machinery, 1993), p. 125.

7. An account of Shannon's theory can be found in Claude E. Shannon and Warren Weaver, *The Mathematical Theory of Communication* (Urbana: University of Illinois Press, 1949).

8. *Cybernetics* (Eighth Conference, 1952), p. 22.

9. Donald M. MacKay, "In Search of Basic Symbols," in *Cybernetics* (Eighth Conference, 1952), pp. 181-221. A fuller account can be found in Donald M. MacKay, *Information, Mechanism, and Meaning* (Cambridge, Mass.: MIT Press, 1969).

10. Heims, *Cybernetics Group* (above, n. 1), p. viii.

11. David B. Parkinson designed a robot gun director for the M-9 antiaircraft gun after having a dream in which he was a member of a Dutch antiaircraft battery that had a marvelous automatic robot gun. The design was implemented in 1942 by Parkinson and two colleagues from Bell Laboratories, Clarence A. Lovell and Bruce T. Weber; it played an important role in the defense of London during World War II. For an account of the gun and a reproduction of Parkinson's original drawing, see the catalogue from the IBM exhibit on the history of the information machine, *A Computer Perspective: A Sequence of 20th-Century Ideas, Events, and Artifacts from the History of the Information Machine,* ed. Glen Fleck (Cambridge, Mass.: Harvard University Press, 1973), pp. 128-129.

12. *Cybernetics* (Seventh Conference, 1951), p. 155.

13. Claude E. Shannon, "Presentation of a Maze-Solving Machine," *Cybernetics* (Eighth Conference, 1952), pp. 173-180.

14. *Cybernetics* (Eighth Conference, 1952), p. xix.

15. Ibid.

16. Ibid., p. 173.

17. W. Ross Ashby, "Homeostasis," in *Cybernetics* (Ninth Conference, 1953), pp. 73-108.

18. Ibid., p. 79.

19. Ibid., p. 73.

20. Ibid., pp. 73-74.

21. Ibid., p. 75.

22. Ibid., p. 107.

23. Ibid., p. 97.

24. Ibid., p.100.

25. John Stroud, "The Psychological Moment in Perception," in *Cybernetics* (Sixth Conference, 1949), pp. 27-63, esp. pp.27-28.

26. *Cybernetics* (Sixth Conference, 1949), p.147.

27. Ibid., p. 153.

28. Ibid.

29. J. Y. Lettvin, H. R. Maturana, W. S. McCulloch, and W. H. Pitts, "What the Frog's Eye Tells the Frog's Brain," *Proceedings of the Institute of Radio Engineers* 47: 11 (1959): 1940-1959.

30. Ibid., p. 1950.

31. H. R. Maturana, G. Uribe, and S. Frenk, "A Biological Theory of Relativistic Color Coding in the Primate Retina," *Archivos de biologia y medicina experimentales*, suppl. 1 (Santiago, Chile: 1969).

32. Heinz von Foerster, "Molecular Ethology: An Immodest Proposal for Semantic Clarification," in *Observing Systems* (above, n. 2), pp. 150-188.

33. Brenda Laurel and Rachel Strickland showed a video about *Placeholder* at the Fourth International Cyberconference at Banff Centre for the Arts in May 1994. The simulation is so complex that it requires ten computers to run it, including two Onyx Reality Machines; it will probably never be shown publically for this reason, much less marketed. It demonstrates how far the industry still is from creating a multiperson, interactive simulation that takes embodiment fully into account.

34. See, for example, Mark Pesce's presentation at the Third International Cyberconference in Austin, Texas, May 1993, entitled "Final Amputation: Pathologic Ontology in Cyberspace" (forthcoming in the electronic journal *Speed*).

35. Howard Rheingold, *Virtual Reality* (New York: Summit Books, 1991).

36. Bruce Mazlish, *The Fourth Discontinuity: The Co-Evolution of Humans and Machines* (New Haven: Yale University Press, 1993).

37. Moravec, *Mind Children* (above, n. 5), pp. 1-5.