THE CALCULUS OF INTRICACY

Can the Complexity of a Forest Be Compared with That of Finnegans Wake?

by SETH LLOYD

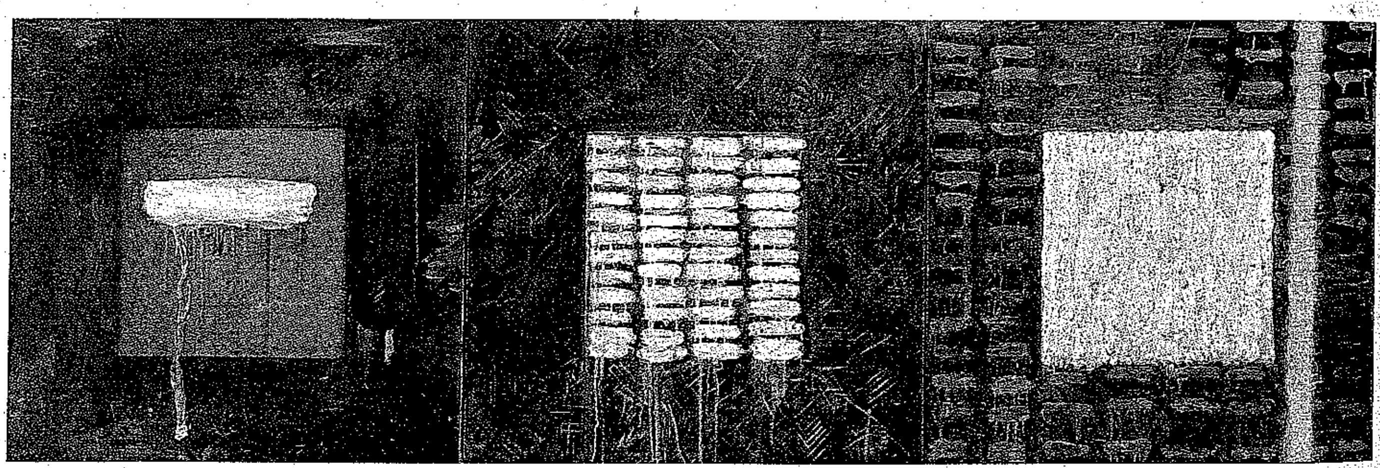

Pat Steir, Beautiful Painting *(detail), 1977–78*

IN THE TOWN OF Oracle, Arizona (population 5,500), just north of Tucson at the edge of the Sonora desert, stands a steel-and-glass structure that looks like a huge commercial greenhouse. Through the four acres of glass that frame this latter-day crystal palace the visitor can make out a lush and diverse stock of plant life. Not so readily apparent, though, are the five years of design, thought and experimentation that inform every detail of the living process expected to unfold and thrive behind its walls.

Inside are seven simulated biomes, or biological settings, including a desert, a savanna, a rain forest, a marshland, an urban settlement, an agricultural district and a twenty-five-foot-deep ocean with artificial wave action to stimulate the growth of a coral reef. Standing in for the wind, rain and other atmospheric forces is a technology that, to the best of the designers' ingenuity, will simulate the corresponding natural processes. Computers will regulate temperature, humidity, tides and the change of seasons. And below the flora and fauna on the floor of the greenhouse is a dense network of machinery designed to cool and filter the air, store water during periodic dry spells and recycle animal waste.

The system is designed to be closed to air, water and all forms of animal, insect and microbial life—except that, like the earth, it will exchange energy and information with its external environment. Beyond those two flows, if the system can survive at all, it will be entirely self-sustaining. In most contexts the word *microcosm* is applied imprecisely, often to make some point about local social dynamics. But in Oracle, Arizona, the expression has a much less figurative connotation, for inside the structure called Biosphere II (the private developers of Biosphere II regard the earth itself as Biosphere I) there is indeed a microcosm, a "little world," a bonsai replica of the terrestrial environment.

Late in the fall, if all goes as planned, eight volunteers, four men and four women, will enter Biosphere II to live for two years. Granted modest quarters and a 24,000-square-foot farm with which to feed themselves, the volunteers will live independently of the outside world. The idea is that the entire enterprise will teem with interconnected life-forms: symbiotic bacteria; plants for recycling carbon dioxide; hummingbirds to pollinate jungle plants; termites to break down savanna grasses; ginger trees to protect the rain forest from excessive sunlight; fishes, fowl, fruits, vegetables and, of course, the eight intrepid members of *Homo sapiens.*

The org
reasons for
that such a
someday e
other inho
implication
all the assi
tions of the
guarantee
What m
predict in c
constructio
a familiar
ness is a co
componen
instance) p
componen
of the bea
spite of ev
simple and

seems to b
complexity.
tems will u
been taken
complicated
Is it, ther
an environ
defined? Ca
what are ot
scale of con
measure fo
earth: comp
ronment, th
complexity
thermodyna
In the p
questions h
fluence of
abstract lo
remains in r
strongly sug
sizes its gen
of a physic
processed in

ACY

ıke?

—except that,
d information
e two flows, if
entirely self-
osm is applied
ut local social
pression has a
the structure
of Biosphere
re is indeed a
ica of the ter-

ht volunteers,
here II to live
nd a 24,000-
elves, the vol-
de world. The
with intercon-
s for recycling
jungle plants;
inger trees to
nlight; fishes,
eight intrepid

The organizers of the project have offered a number of reasons for going to all the trouble. Some have suggested that such an enclosed, self-sufficient environment might someday enable people to inhabit the moon, Mars or other inhospitable settings. But for all the millennial implications of a working, self-sustaining biosphere, for all the assiduous planning, for all the sanguine expectations of the project developers, there is by no means any guarantee that the two-year experiment will succeed.

What might cause the ship to founder? No one can predict in detail, of course, but the years of planning and construction at Biosphere II have repeatedly driven home a familiar ecological lesson: environmental connectedness is a complex business. An adjustment made to any component of the system (the height of the waves, for instance) propagates its effects onto virtually every other component (the relative humidity of the air, the stability of the beach), requiring adjustments there as well. In spite of every attempt to introduce components that are simple and robust, the dynamic equilibrium of the system eems to be purchased with an ever growing filigree of omplexity. If things still manage to go awry, the postmorems will undoubtedly point to measures that, had they een taken, would have made the system even more omplicated than it is.

Is it, then, complexity itself that in some sense sustains n environment? If so, can that complexity be precisely efined? Can one measure it? Can one thereby compare hat are otherwise wildly disparate entities on a single cale of complexity? What are the implications of such a neasure for an understanding of complex systems on arth: computers, brains, insect societies, the global environment, the origin of life? Could it be that in some sense omplexity is the firebreak that retards the inexorable hermodynamic dissolution of the world?

In the past forty years an attempt to answer these uestions has evolved into a subtle study at the confluence of physics, mathematics, computer science and bstract logic. Although the concept of complexity emains in many ways elusive, my work and that of others rongly suggests an approach to the concept that emphazes its genesis: any universal measure of the complexity f a physical system must account for the information rocessed in the course of creating that system, whether by natural selection or computation, by bugs or human beings.

TO THE EXTENT THAT Biosphere II succeeds, it will be in virtual defiance of a fundamental principle of physics, the second law of thermodynamics. Formulated in the middle of the nineteenth century by the physicists Rudolf J. E. Clausius of Germany and Ludwig E. Boltzmann of Austria, among others, the second law guarantees that, to pinch a phrase from W. B. Yeats, things fall apart. That is, left on their own, all ordered systems—a hot jet of steam, a human body, a political hierarchy, even a Toyota—eventually decay into maximum disorder, and entropy (the amount of randomness or disorder in the universe) tends to increase.

The high hopes for Biosphere II, the gloomy implications of the second law notwithstanding, rest on the confidence that it can adequately model the earth's capacity for self-organization and renewal. That capacity, founded on an intricate web of nuance spun by four billion years of evolution, seemingly has managed to hold the second law at bay. From the moment a few scattered primordial molecules banded together to form the first bacterium, life on this planet has shown a continual trend toward a highly ordered but labyrinthine complexity. The living things around us have a long history of finding new ways of fueling themselves with energy and of extracting work and advantage from the decay of that energy.

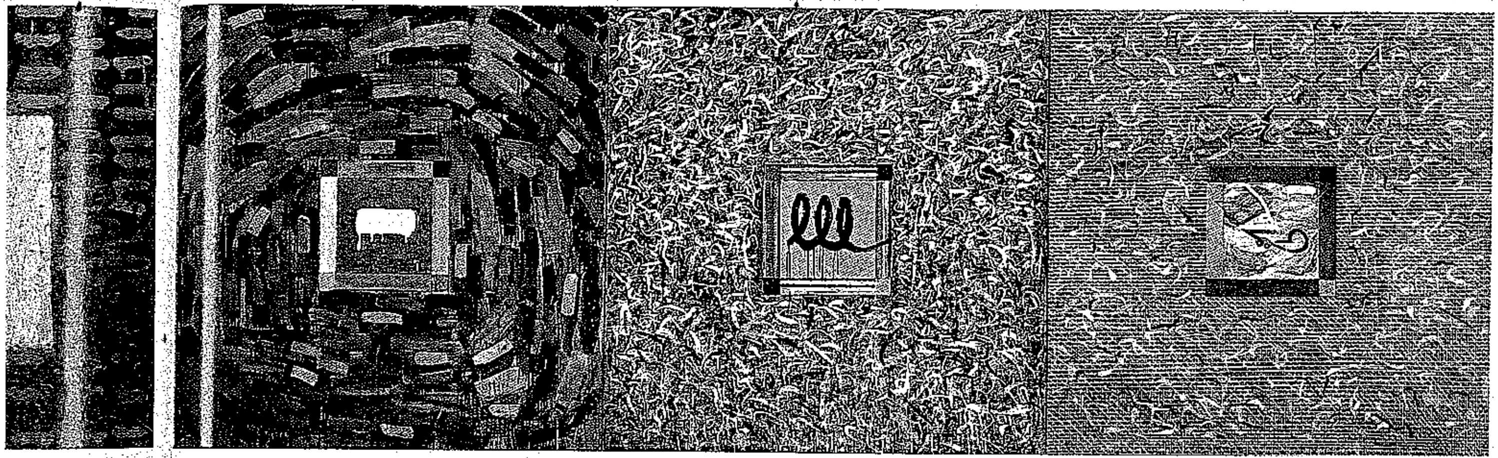

Pat Steir, Beautiful Italian Painting, *1980*

Bacteria, for instance, have tried out more than a trillion trillion genetically programmed strategies for survival and found them either workable or wanting, worth keeping or disposable. Natural selection has made the resultant organisms subtle and highly developed: bacteria today exhibit a range of talents, from metabolizing sulfur to fooling the human immune system, that allow them to exploit a great variety of energy resources. This hard-won flexibility has resulted from the constant intake of information, received from myriad environmental stimuli over the millennia and converted into layer upon layer of sophisticated potential for environmental response.

The growth of complexity and order in many systems is empirically obvious, and it seems conceptually to parallel the growth of disorder guaranteed by the second law of thermodynamics. Yet it has acquired no widely accepted scientific form. One apparently crucial element in any reasonable measure of complexity is the information processed or exchanged by the system under study.

Can that information be measured? It can—but only if a precise definition of information is available. Such a definition was provided in 1948 by the mathematician Claude E. Shannon, then at Bell Telephone Laboratories, in a paper titled "The Mathematical Theory of Communication." In Shannon's theory information is conveyed by making a series of choices from among a set of alternatives. In other words, where there is possibility, there is information. Often the set of alternatives boils down to a yes-or-no proposition: some kind of prearranged mark—a dot, a surge in voltage, a magnetized region—either is present or it is not. The two alternatives are represented by the two binary digits 1 and 0, and so the bit (short for binary digit) has become the basic unit of recording information.

A familiar example of this principle is the game of twenty questions. In a series of yes-or-no queries ("Is it bigger than a bread box?") the contestant tries to identify some mysterious person, place or thing. According to the simplest measure of information content, each yes or no answer conveys a single bit of information. A more sophisticated measure recognizes that some answers are more revealing than others and rewards correct answers with more information. If the first of your twenty questions is, "Was the person more than five feet eight inches tall?" a yes or a no will give you roughly the same amount of information. But if you ask, "Was this person a dramatist who, legend has it, was killed by a tortoise dropped from the beak of an eagle?" there is a much wider disparity. A no will teach you precious little, whereas an unlikely yes should win you the game—assuming that, since you knew enough to ask, you know enough to realize the mystery person is Aeschylus.

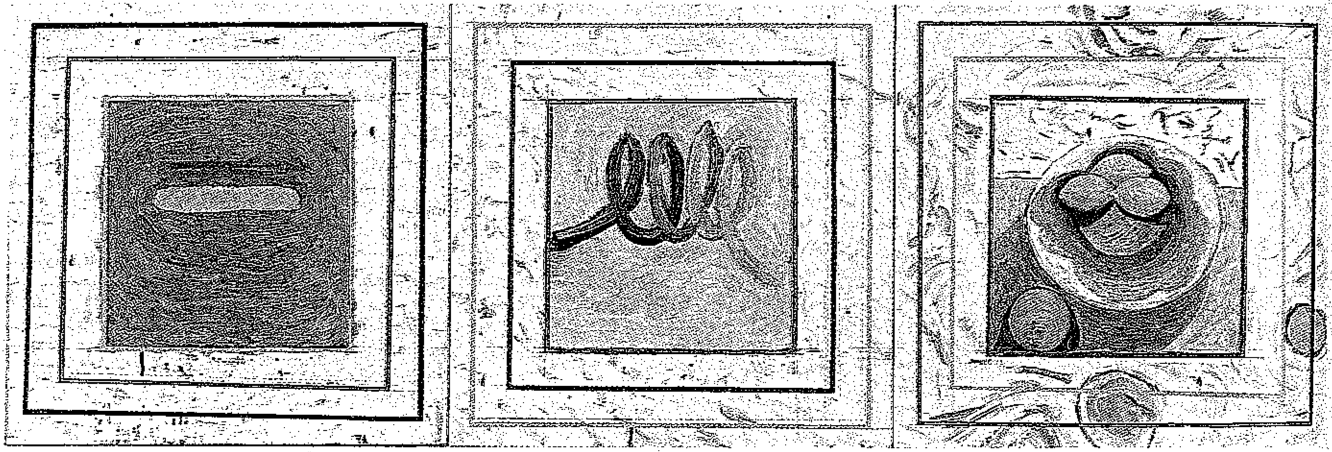

Pat Steir, Triptych with Still Life, *1981*

THUS THE MORE UNLIKELY the alternative, the more information you stand to gain from it. Information is measured on a kind of Richter scale for surprise: for every factor of two in improbability, one more bit of information is conveyed. (As a consequence, the way to get the most information in a game of twenty questions is to ask questions for which a yes and a no are equally probable. When the answers have different probabilities, the answer you are more likely to get conveys less than one bit of information.)

To understand what information theory might say about complexity, it is useful to explore the close relation between information theory and the second law of thermodynamics. Entropy and information are complementary quantities. The entropy of the helium in a balloon measures the number of ways the positions and velocities of the helium molecules could be arranged without changing the temperature, pressure or shape of the balloon. In other words, entropy measures how much information is lacking about the configuration of the molecules in the balloon, given its macroscopic description. Entropy is a measure not only of disorder but of ignorance.

Conversely, the more orderly a system is, the more redundancy it carries, and thus the less information is needed to describe it. For example, suppose that one constructs a sequence of letters by selecting a letter at random and then writing it eight times: bbbbbbbb. One can hardly ask for a more orderly arrangement than that, and its orderliness reflects its low information content. In contrast, a string of letters constructed by choosing letters at random—say qxkeufhjdoewjnatybm—embodies the ultimate in disorder: there is no rhyme, reason or redundancy to it, and it represents a large amount of information; each letter is a complete surprise.

ONE MIGHT GUESS THAT the complexity of a system can be measured by the amount of information needed to describe it. Returning to the string of letters, this definition seems to make perfect sense: qxkeufhjdoewjnatybm appears, by a comfortable margin, more complex than bbbbbbbb. Yet were this approach taken as an accurate gauge of complexity, it would turn literary scholarship on its ear. By this measure a chimpanzee typing 650 pages of random alphanumeric characters would in short order produce a work not only as long as but far more complex than James Joyce's *Finnegans Wake,* which was completed only after seventeen years of diligent labor. English, Gaelic, Latin and any of the other languages Joyce employed are redundant: not all combinations of letters or words are allowed, and though *Finnegans Wake* is full of surprises, *q*'s are still followed, more

ferent prob-
get conveys

might say
ose relation
law of ther-
omplemen-
in a balloon
ld velocities
ed without
: of the bal-
much infor-
e molecules
on. Entropy
ance.
, the more
ormation is
hat one con-
er at random
). One can

an that, and
content. In
osing letters
nbodies the
n or redun-
of informa-

plexity of a
nount of in-
the string of
fect sense:
ible margin,
is approach
would turn
ure a chim-
neric charac-
only as long
's *Finnegans*
een years of
of the other
t all combi-
though *Fin-*
owed, more

often than not, by *u*'s. By Shannon's definition, then, the information content of the book is less than that of the chimpanzee's magnum opus, in which each letter is fresh, new, untainted by scholarship and completely random. Complexity is not to be equated with information content alone. (It is interesting to note in this context that *Finnegans Wake* probably has a higher information content on each page and lower redundancy overall than any other novel. It is this very lack of predictability that makes it seem largely random and difficult to read.)

Information fails as a measure of complexity for physical systems as well. A gas, for instance, displays its simplest properties at equilibrium. But it is exactly at equilibrium—the state of maximum entropy—that the configuration of the molecules in the gas demands the most information to describe.

AN ALTERNATIVE APPROACH to measuring complexity, one that amounts to a computational analogue of Shannon's information content, was developed independently in the 1960s by Raymond J. Solomonoff, then of the Zator Company in Cambridge, Massachusetts; Gregory J. Chaitin, who was then an undergraduate at the City College of the City University of New York; and Andrei Kolmogorov of the Mathematics Institute in Moscow. In this approach complexity is measured by algorithmic information content—the size of the shortest algorithm, or computer program, that can generate the number, object or state under scrutiny. For example, the long decimal expansion .285714285714285714285714285714285714 can be generated by a short algorithm: *print 2/7 to 36 digits*. On the other hand, it is hard to see how a number such as .8329513 could be generated by an algorithm shorter than the number itself. Its apparent randomness suggests a handy definition: a sequence of numbers is random if it cannot be algorithmically compressed.

To apply the concept of algorithmic information content to the natural world, imagine reducing the description of any physical system to a sequence of numbers, displayed in binary form as a string of ones and zeros. For example, the molecular configuration of a crystal might become 10101010101010101010..., a regular binary array that can be described with a brief algorithm (*print 10* n *times*). But a description of the random behavior of gas molecules at equilibrium, converted to an arbitrary sequence of ones and zeros, cannot be abbreviated, any more than can the chimpanzee's stochastic jottings.

Although it is a useful measure of randomness, algorithmic information content also fails as a measure of complexity. Like Shannon's information content, it would assign less complexity to *Finnegans Wake* (which many regard as the most impenetrable novel in print) than it would to a random string the same length as Joyce's text. It would confer a high average complexity to configurations of molecules in thermal equilibrium. Finally, since it depends on the language in which the working algorithm is expressed, it cannot be an entirely unambiguous quantity. After all, one can say some things more concisely in English than in Japanese. Because different computers adopt different machine languages, algorithmic information content depends on specific properties of the computer on which the algorithm is to run.

Thus both Shannon's information content and algorithmic information content fail to capture our intuitive understanding of the concept of complexity, either by assigning a higher complexity to the completely random but simple than to the ordered but intricate or by assigning a high complexity to such straightforward physical situations as thermal equilibrium. A computational measure that values intricacy above randomness and simplicity does exist, however. The measure was formulated by Charles H. Bennett, a physicist at the IBM Thomas J. Watson Research Center in Yorktown Heights, New York; it is called logical depth.

BENNETT'S APPROACH appeals to the intuitive idea that a complex object is difficult to construct. To measure the complexity of an object—a number, say—he begins by determining the most plausible computer program that would give the number as output. The program would be selected from among the shortest programs that can construct the number on a given computer. That choice made, Bennett then defines logical depth as the number of logical steps (on the computer he would count machine cycles) that connect the most plausible program with the constructed number.

In this scheme, for example, random numbers are logically shallow: the shortest program that can generate a random number *X* is written *print* X, and takes only a short time to run. And straightforward, ordered numbers are not logically deep either: the number made up of a billion ones is generated by issuing the command *print 1 a billion times*, which also runs fast. In contrast, there are logically deep questions, for instance, in number theory or in chess: What is the seventy-ninth term in the series of so-called Fermat prime numbers? Does white have a winning strategy after the gambit, queen's bishop declined? True, in each case a short program can supply the answer—but only after a long time and a great many machine cycles. Something as organizationally complex as a living creature, meanwhile, would call for an elaborate program running for many years in order to reflect the complex sequence of events in the creature's evolution.

Bennett's measure would make for a fairly accurate comparison between the relative complexities of *Finnegans Wake* and the output of the typing chimpanzee. The complexity of Joyce's novel could be measured by adding the number of computer cycles it would take to simulate the circumstances that led Joyce to write the novel—how he arrived at each meticulously chosen word, each maddeningly obscure allusion, each act of linguistic innovation. Perhaps one ought to include every biographical detail of Joyce's life that might have had a bearing on his labors—his writing of *Ulysses* and other earlier works, his education and childhood in Dublin, his expatriation, his endless series of eye operations, his relationships with his family and with anyone else who crossed his path. One might even go so far as to count up the information processed in all the evolutionary steps leading to the birth of a son to John Stanislaus and Mary Jane Murray Joyce on February 2, 1882. Although one could similarly trace the origins of the chimpanzee, it is a safe bet that Joyce and his tortuous classic would come off as being appreciably more complex.

But Bennett's construct is, by his own admission, impractical. The first step in arriving at logical depth—finding the most plausible algorithm for whatever you are measuring—is highly problematic: There is no systematic way of determining such an algorithm. Furthermore—for reasons closely related to the famous theorem of the logician Kurt Gödel about the undecidability of certain mathematical statements—you can never be absolutely certain, even if you have guessed or otherwise divined a clear and simple algorithm, that what you have found is one of the shortest programs for constructing the object in question. Another difficulty is that, as with algorithmic information content, the logical depth of a number depends on the specific computer on which the most plausible algorithm is to run. For example, quantum-mechanical calculations of the strong nuclear forces between protons, neutrons and quarks require far fewer machine cycles on a massively parallel processing machine than they do on a conventional computer.

SUCH FLAWS NOTWITHSTANDING, it is an appealing idea to identify the complexity of a thing with the amount of information processed in the most plausible method of its creation. The problems with logical depth derive largely from its dependence on a specific computer. Information, however, is the currency not only of computation but also of all of nature. Light scattering from a surface gets and transforms information about the shape and texture of the surface; natural selection gathers information about the environment and summarizes it in a species' genes.

With this in mind my late colleague the physicist Heinz R. Pagels and I proposed a method of assigning a number to the total amount of information processed during the evolution of a physical or a biological system. Identify the most plausible method of constructing a thing not with the most concisely encoded program for creating its digital representation but with the most plausible scientifically determined sequence of events that lead to the thing itself. Measure complexity not by counting the number of computer cycles required by the computation of that representation but by reckoning the total amount of thermodynamic and informational resources required by the physical construction process. Quantifying the complexity of the system in terms that relate to its thermodynamics opens up the possibility of pinpointing physical constraints on the system that help or hinder the development of complexity.

By identifying the complexity of a thing with the amount of informational and thermodynamic effort involved in putting it together, one avoids ascribing high complexity to things such as random sequences of letters that carry much information but are easy to assemble. A living creature such as a Brahma bull is extremely complex; the amount of information processed in four billion years of evolution is immense. Two Brahma bulls, however, are only slightly more complex than a single bull: given one bull, all that is required to produce another is a complementary cow.

To take an example in which the amount of information processed can be gauged more accurately, consider a culture of bacteria. The gene pool of each generation in the culture is the same as that of the preceding generation, except for differences in the frequency of certain genes and mutations in the structure of those genes. The culture evolves by natural selection: a gene achieves greater representation, or a mutation survives, according to whether the individual bacteria carrying the gene or mutation cope successfully with their environment and are able to reproduce. Over the generations the species absorbs and processes information from its surroundings as natural selection "computes" the genetic configurations that are most adaptable to the slowly changing environment. But there is no computer here. How does one measure the amount of information processed?

To begin, find the length of the message needed to label all possible genetic configurations of the subsequent generation—that is, the location and type of all possible mutations and all possible changes in the frequency of different genes—given the composition of the gene pool of the present generation. That length is what Pagels and I called the thermodynamic depth of the reproductive process over one generation. It is simply the informational capacity the species brings to bear in regeneration, the total amount of "space" available to absorb information from the environment.

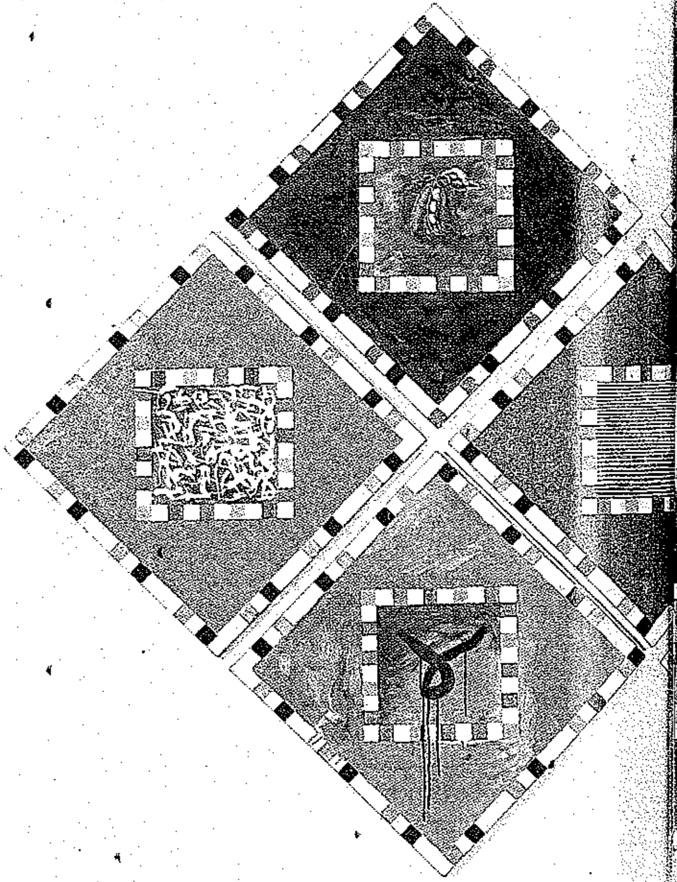

Pat Steir, Being/Been, *1980*

IT WOULD BE highly unusual for a species to spend its entire information-processing capacity in one generation. To understand what information is used, imagine adding a twist of lemon to the culture. The bacteria swim around as before, but genes and mutations that enhance bacterial reproduction in the acidic environ-

enes. The cul-
ıchieves greater
·, according to
ıg the gene or
ıvironment and
ons the species
ts surroundings
ıetic configura-
changing envi-
How does one
essed?
sage needed to
the subsequent
e of all possible

requency of dif-
he gene pool of
hat Pagels and I
eproductive pro-
ie informational
generation, the
orb information

species to spend
capacity in one
nation is used,
ie culture. The
s and mutations
e acidic environ-

---

ment will occur with increasing frequency in subsequent generations. There are always a certain number of genes and mutations, however, that have no bearing on the reproductive success of the culture. When the length of the message needed to specify such residual genetic variation is subtracted from the overall information-processing apacity, or thermodynamic depth, the difference is a quantity Pagels and I called breadth—a measure of the information capacity of the species that is put to work.

A culture of bacteria kept at its favorite temperature and given an unlimited supply of food undergoes genetic variation largely unconstrained by the environment. Such conditions generate little breadth; the bacteria process virtually no new information about their environment. But

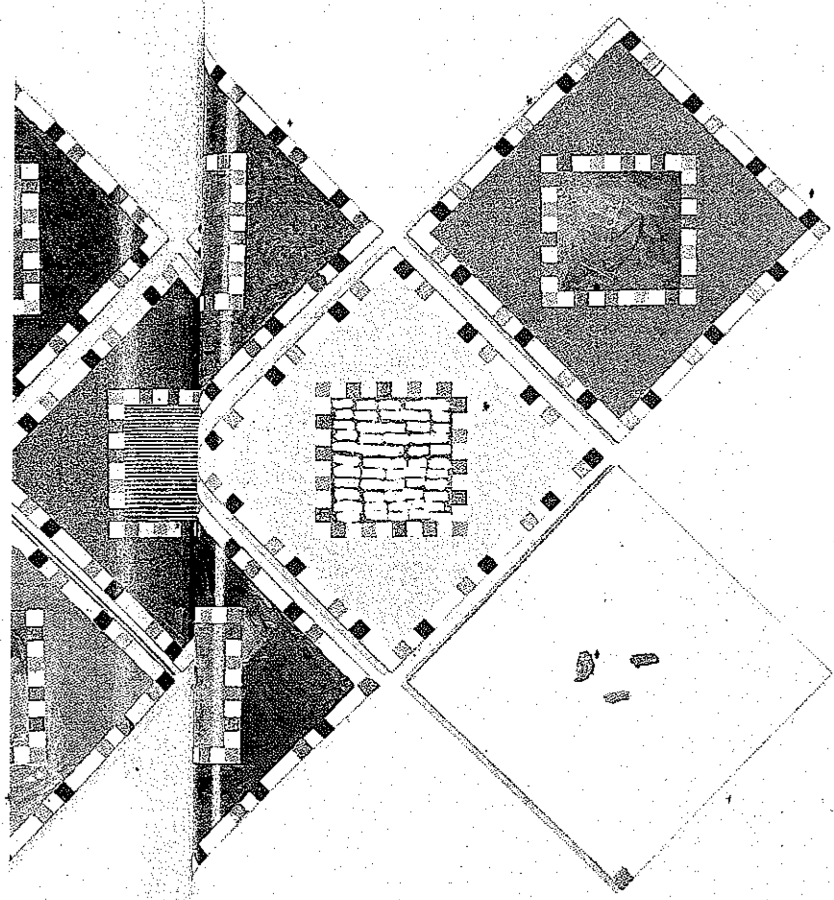

he same bacterial culture, starved and subjected to exeme variations of temperature, must deploy much more f its genetic informational capacity in order to survive.

The rate at which a bacterial culture processes informaon can be prodigious. Imagine first a solitary specimen of *scherichia coli*. In two hours or so the bacterium can act in ne of three ways: it can survive and multiply, survive ithout reproducing, or die. In a stable population those ehaviors are about equally likely over the given period. hus the bacterium can register, or take in, one of three ternatives—fewer than two bits of information—in the vo-hour span.

But wait: If the bacterium lives and reproduces, one ust also take into account the possible changes in its netic sequence and that of its immediate offspring (through, say, mutation, recombination or transcription error during reproduction). In that case the number of potential outcomes jumps dramatically. Instead of just three possible states, there are now at least several billion. For example, the bacterium might survive after a cosmic ray knocks out base pair number 55 in its DNA sequence; or the bacterium might reproduce, but in the process make a transcription error in which cytosine is substituted for guanine in the 50,023rd base pair of the offspring's DNA—and so on. The amount of information needed to describe the location and type of all such mutations and transcriptions is at least several hundred bits.

Thus each bacterium in a culture undergoing natural selection can register as many as a few hundred bits of information in its life span. (If a bacterium dies or does not reproduce, it registers only a few bits of information for the succeeding generation.) Assuming the bacteria do not exchange genetic information with one another, the total information needed to specify the behavior and genetic variation of the entire culture is the simple sum of the corresponding information for the individual bacteria. If the population of the culture is a billion *E. coli*, its thermodynamic depth is roughly 100 billion bits a generation. That translates into an information-processing capacity of several million instructions a second.

The amount of that capacity put to use in a given period—the culture's breadth—depends on the stress imposed by the surroundings. If the effect of the environment on bacterial survival and reproduction is small, and if most genetic changes have little effect on the bacteria's ability to survive or reproduce in the environment, the culture will process information much more slowly than this estimate. But if the environmental stress is high, and if many genetic variations have an effect on the viability of the species, the full adaptive power of the bacteria is brought into play. The culture can apply what amounts to the power of a mainframe computer—not to problems such as long division but to the problem of survival.

ONE DIFFICULTY with the two computationally based measures of complexity I discussed earlier—algorithmic complexity and logical depth—is that each is subject to conceptual limitations that are imposed by Gödel's undecidability theorem. These limitations cannot readily be eliminated; instead they arise from a deep and, to some, disturbing fact about the world. Gödel's announcement of his theorem in 1931 sent a temblor that passed straight through the small community of logicians and metamathematicians and on to a much wider audience of mathematicians and even, to some extent, the lay public; the aftershocks continue to be felt in what seem to be ever widening contexts.

In a rigorous proof, Gödel showed that any formal system of logic elaborate enough to include the axioms of arithmetic must either be internally inconsistent or else include true propositions that can be neither proved nor disproved from the axioms and rules of inference that govern the system. In other words, arithmetic is incomplete.

Now remember that to measure algorithmic complexity you have to determine the shortest string of bits capable of specifying the object in question. As Gregory Chaitin, now at the IBM Thomas J. Watson Research Center, pointed out fifteen years ago, Gödel's theorem makes it impossible to know that the shortest string has been encountered. Logical depth is subject to a similar constraint, since it requires knowledge of the shortest possible algorithm for specifying the complex object.

To appreciate the difficulty, consider a problem familiar to puzzle enthusiasts: What is the largest number that can be formed out of three 4s? (The standard mathematical notation can be included at no cost to the count.) Clearly $4\times4\times4$ (64) does better than $4+4+4$ (12), but 444 is better still. A tower of powers, 4^{4^4}, however, wins the day: the number is greater than 10^{154}. Now ask essentially the same question after you allow both the number of symbols and the length of the string to grow: for example, What is the smallest number that requires more than sixty-eight symbols to be specified? But the phrase "the smallest number that requires more than sixty-eight symbols to be specified" is only sixty-eight characters long (not counting spaces and quotation marks). Hence the phrase seems to specify a number that cannot be specified in a phrase of that length. The paradox of such examples leads to a striking conclusion: one cannot prove the truth of any answer to such a question, for if one could, a contradiction would ensue. The only way to avoid contradiction is to admit that certain questions are undecidable.

BREADTH, the measure of complexity Pagels and I devised, is free of such logical constraints, provided one calculates it with experimentally determined probabilities. Nevertheless, to the extent that it relies on extrapolation beyond experiment, breadth may suffer from the same logical paradox that infects all of science: as Solomonoff has pointed out, the most plausible scientifically determined sequence of events for constructing a particular object is analogous to the shortest algorithm for specifying the object. Such a sequence may therefore be impossible to determine.

The astrophysicist Wojciech H. Zurek of the Los Alamos National Laboratory suggested last year, in an article published in *Nature,* that the logical paradoxes of algorithmic complexity theory have implications for the hard physical world of complex machines. To see how Gödel's theorem can have an effect on the efficiency of an engine that processes complicated information, consider the contemporary automobile. Many cars today are equipped with sensors in the exhaust pipe, at the springs and at the accelerator cable that transmit information to microprocessors. In turn, the microprocessors apply the information to regulate the ratio of fuel to air in the carburetor, to adjust the suspension of the car on the road, even to change the timing of the spark. The microprocessors encode the information they receive in binary form, process it and send out their instructions. When the information is of no further use, or when the information storage capacity of the microprocessor is filled, old information is erased. But the more information one discards, the greater the increase in entropy. Thus to operate at maximum thermodynamic efficiency, the car microprocessors must encode their information as concisely as possible. But as I have noted above, the most concise form to which a piece of information can be compressed cannot itself be computed. No matter how hard they try, the microprocessors will always erase more information than necessary.

The amount of thermodynamic inefficiency imposed on an automobile by Gödel's theorem is infinitesimal, considerably less than a billionth of a billionth of a billionth of the amount of waste avoided by changing the oil regularly. Still, the idea that logical paradox leads to wasted resources—a loss now known as logical friction—is compelling: no computation that must take the maximum possible advantage of its memory resources, no scientific theory that represents itself as the simplest (shortest) possible explanation for a process, can avoid being caught in the net of Gödel's remarkable result.

THAT ALL OF SCIENCE IS, in principle, subject to the restrictions of Gödel's theorem does not prevent the legitimate application of our new measures for comparing the amount of complexity generated by two radically different processes. A sand castle is thermodynamically deeper than a sand dune, because the conscious act of a child's packing wet sand into a hard mass and then sculpting it into the facsimile of a turreted medieval fortress involves more information than do the geological and meteorological forces that conspire to make dunes. For one thing, when measuring the breadth and thermodynamic depth of the sand castle one must account not only for the natural processes that result in a sandy beach, but also for the evolutionary processes that created the child. (The approach here is similar to that of applying Bennett's logical depth to *Finnegans Wake*—except that one is no longer constrained by machinery. In practice, of course, there are serious stumbling blocks, since no one has the slightest idea what, say, the breadth of a child might be. But that is no objection to the principle behind the definition.)

Even more intriguing, one can speculate about the rate at which complexity is evolving in people. In our species complexity arises not from the shuffling of the genes alone. In fact, by insulating us from the caprices of the environment, our technology, culture and ethical principles slow the rate at which humanity is growing more complex—just as a protective shield would diminish the breadth of a culture of bacteria. But the decrease in the growth rate of the complexity of the human genome is well compensated for by the capacity of new technologies to process and exchange information on our behalf.

The questions, How does complexity arise? When does it flourish? Under what conditions does it decay? are not academic. The study of complexity aims at understanding not only the origins of life but also the possibilities for its future evolution. Gauging the complexity of an ecosystem (or, indeed, the biosphere) might help in assessing the implications of tropical deforestation or the greenhouse effect—whether, for instance, it is the enormously complicated tangle of interacting life-forms in the Brazilian rain forest that makes the survival of that region a more delicate matter than it is in, say, Yosemite National Park. Gauging the complexity of the immune system might help determine, economically, to what degree and in what directions we as a species need to develop our own complexity to meet the challenges of disease. For, after all, a species stumped by an intractable problem does not merely cease to compute. It ceases to exist. ●

SETH LLOYD is a postdoctoral fellow in the physics, mathematics and astronomy division of the California Institute of Technology.

LAND FILL
A Political Ec
by Edwin N.
University of
17.95 paperb

n 1962 my
hilippines, i
he modern w
he most "pr
legritos knov
n the rain for
vore little G-s
vith bows and
elieved we w
ving much as
he beginning
he utter isola
he end of the
Two weeks
ves," howev
utside our hu