

Gödel's Theorem and Information

Gregory J. Chaitin

IBM Research, P.O. Box 218, Yorktown Heights, New York 10598

Received April 14, 1982

Gödel's theorem may be demonstrated using arguments having an information-theoretic flavor. In such an approach it is possible to argue that if a theorem contains more information than a given set of axioms, then it is impossible for the theorem to be derived from the axioms. In contrast with the traditional proof based on the paradox of the liar, this new viewpoint suggests that the incompleteness phenomenon discovered by Gödel is natural and widespread rather than pathological and unusual.

1. INTRODUCTION

To set the stage, let us listen to Hermann Weyl (1946), as quoted by Eric Temple Bell (1951):

We are less certain than ever about the ultimate foundations of (logic and) mathematics. Like everybody and everything in the world today, we have our "crisis." We have had it for nearly fifty years. Outwardly it does not seem to hamper our daily work, and yet I for one confess that it has had a considerable practical influence on my mathematical life: it directed my interests to fields I considered relatively "safe," and has been a constant drain on the enthusiasm and determination with which I pursued my research work. This experience is probably shared by other mathematicians who are not indifferent to what their scientific endeavors mean in the context of man's whole caring and knowing, suffering and creative existence in the world.

And these are the words of John von Neumann (1963):

...there have been within the experience of people now living at least three serious crises... There have been two such crises in physics—namely, the conceptual soul-searching connected with the discovery of relativity and the conceptual diffi-

culties connected with discoveries in quantum theory... The third crisis was in mathematics. It was a very serious conceptual crisis, dealing with rigor and the proper way to carry out a correct mathematical proof. In view of earlier notions of the absolute rigor of mathematics, it is surprising that such a thing could have happened, and even more surprising that it could have happened in these latter days when miracles are not supposed to take place. Yet it did happen.

At the time of its discovery, Kurt Gödel's incompleteness theorem was a great shock and caused much uncertainty and depression among mathematicians sensitive to foundational issues, since it seemed to pull the rug out from under mathematical certainty, objectivity, and rigor. Also, its proof was considered to be extremely difficult and recondite. With the passage of time the situation has been reversed. A great many different proofs of Gödel's theorem are now known, and the result is now considered easy to prove and almost obvious: It is equivalent to the unsolvability of the halting problem, or alternatively to the assertion that there is an r.e. (recursively enumerable) set that is not recursive. And it has had no lasting impact on the daily lives of mathematicians or on their working habits; no one loses sleep over it any more.

Gödel's original proof constructed a paradoxical assertion that is true but not provable within the usual formalizations of number theory. In contrast I would like to measure the power of a set of axioms and rules of inference. I would like to be able to say that if one has ten pounds of axioms and a twenty-pound theorem, then that theorem cannot be derived from those axioms. And I will argue that this approach to Gödel's theorem does suggest a change in the daily habits of mathematicians, and that Gödel's theorem cannot be shrugged away.

To be more specific, I will apply the viewpoint of thermodynamics and statistical mechanics to Gödel's theorem, and will use such concepts as probability, randomness, entropy, and information to study the incompleteness phenomenon and to attempt to evaluate how widespread it is. On the basis of this analysis, I will suggest that mathematics is perhaps more akin to physics than mathematicians have been willing to admit, and that perhaps a more flexible attitude with respect to adopting new axioms and methods of reasoning is the proper response to Gödel's theorem. Probabilistic proofs of primality via sampling (Chaitin and Schwartz, 1978) also suggest that the sources of mathematical truth are wider than usually thought. Perhaps number theory should be pursued more openly in the spirit of experimental science (Polya, 1959)!

I am indebted to John McCarthy and especially to Jacob Schwartz for making me realize that Gödel's theorem is not an obstacle to a practical AI

(artificial intelligence) system based on formal logic. Such an AI would take the form of an intelligent proof checker. Gottfried Wilhelm Leibnitz and David Hilbert's dream that disputes could be settled with the words "Gentlemen, let us compute!" and that mathematics could be formalized, should still be a topic for active research. Even though mathematicians and logicians have erroneously dropped this train of thought dissuaded by Gödel's theorem, great advances have in fact been made "covertly," under the banner of computer science, LISP, and AI (Cole et al., 1981; Dewar et al., 1981; Levin, 1974; Wilf, 1982).

To speak in metaphors from Douglas Hofstadter (1979), we shall now stroll through an art gallery of proofs of Gödel's theorem, to the tune of Moussorgsky's pictures at an exhibition! Let us start with some traditional proofs (Davis, 1978; Hofstadter, 1979; Levin, 1974; Post, 1965).

2. TRADITIONAL PROOFS OF GÖDEL'S THEOREM

Gödel's original proof of the incompleteness theorem is based on the paradox of the liar: "This statement is false." He obtains a theorem instead of a paradox by changing this to: "This statement is unprovable." If this assertion is unprovable, then it is true, and the formalization of number theory in question is incomplete. If this assertion is provable, then it is false, and the formalization of number theory is inconsistent. The original proof was quite intricate, much like a long program in machine language. The famous technique of Gödel numbering statements was but one of the many ingenious ideas brought to bear by Gödel to construct a number-theoretic assertion which says of itself that it is unprovable.

Gödel's original proof applies to a particular formalization of number theory, and was to be followed by a paper showing that the same methods applied to a much broader class of formal axiomatic systems. The modern approach in fact applies to all formal axiomatic systems, a concept which could not even be defined when Gödel wrote his original paper, owing to the lack of mathematical definition of effective procedure or computer algorithm. After Alan Turing succeeded in defining effective procedure by inventing a simple idealized computer now called the Turing machine (also done independently by Emil Post), it became possible to proceed in a more general fashion.

Hilbert's key requirement for a formal mathematical system was that there be an objective criterion for deciding if a proof written in the language of the system is valid or not. In other words, there must be an algorithm, a computer program, a Turing machine, for checking proofs. And the compact modern definition of formal axiomatic system as a recursively enumerable set of assertions is an immediate consequence if one uses the so-called

British Museum algorithm. One applies the proof checker in turn to all possible proofs, and prints all the theorems, which of course would actually take astronomical amounts of time. By the way, in practice LISP is a very convenient programming language in which to write a simple proof checker (Levin, 1974).

Turing showed that the halting problem is unsolvable, that is, that there is no effective procedure or algorithm for deciding whether or not a program ever halts. Armed with the general definition of a formal axiomatic system as an r.e. set of assertions in a formal language, one can immediately deduce a version of Gödel's incompleteness theorem from Turing's theorem. I will sketch three different proofs of the unsolvability of the halting problem in a moment; first let me derive Gödel's theorem from it. The reasoning is simply that if it were always possible to prove whether or not particular programs halt, since the set of theorems is r.e., one could use this to solve the halting problem for any particular program by enumerating all theorems until the matter is settled. But this contradicts the unsolvability of the halting problem.

Here come three proofs that the halting problem is unsolvable. One proof considers that function $F(N)$ defined to be either one more than the value of the N th computable function applied to the natural number N , or zero if this value is undefined because the N th computer program does not halt on input N . F cannot be a computable function, for if program N calculated it, then one would have $F(N) = F(N) + 1$, which is impossible. But the only way that F can fail to be computable is because one cannot decide if the N th program ever halts when given input N .

The proof I have just given is of course a variant of the diagonal method which Georg Cantor used to show that the real numbers are more numerous than the natural numbers (Courant and Robbins, 1941). Something much closer to Cantor's original technique can also be used to prove Turing's theorem. The argument runs along the lines of Bertrand Russell's paradox (Russell, 1967) of the set of all things that are not members of themselves. Consider programs for enumerating sets of natural numbers, and number these computer programs. Define a set of natural numbers consisting of the numbers of all programs which do not include their own number in their output set. This set of natural numbers cannot be recursively enumerable, for if it were listed by computer program N , one arrives at Russell's paradox of the barber in a small town who shaves all those and only those who do not shave themselves, and can neither shave himself nor avoid doing so. But the only way that this set can fail to be recursively enumerable is if it is impossible to decide whether or not a program ever outputs a specific natural number, and this is a variant of the halting problem.

For yet another proof of the unsolvability of the halting problem, consider programs which take no input and which either produce a single natural number as output or loop forever without ever producing an output. Think of these programs as being written in binary notation, instead of as natural numbers as before. I now define a so-called Busy Beaver function: BB of N is the largest natural number output by any program less than N bits in size. The original Busy Beaver function measured program size in terms of the number of states in a Turing machine instead of using the more correct information-theoretic measure, bits. It is easy to see that BB of N grows more quickly than any computable function, and is therefore not computable, which as before implies that the halting problem is unsolvable.

In a beautiful and easy to understand paper Post (1965) gave versions of Gödel's theorem based on his concepts of simple and creative r.e. sets. And he formulated the modern abstract form of Gödel's theorem, which is like a Japanese haiku: there is an r.e. set of natural numbers that is not recursive. This set has the property that there are programs for printing all the members of the set in some order, but not in ascending order. One can eventually realize that a natural number is a member of the set, but there is no algorithm for deciding if a given number is in the set or not. The set is r.e. but its complement is not. In fact, the set of (numbers of) halting programs is such a set. Now consider a particular formal axiomatic system in which one can talk about natural numbers and computer programs and such, and let X be any r.e. set whose complement is not r.e. It follows immediately that not all true assertions of the form "the natural number N is not in the set X " are theorems in the formal axiomatic system. In fact, if X is what Post called a simple r.e. set, then only finitely many of these assertions can be theorems.

These traditional proofs of Gödel's incompleteness theorem show that formal axiomatic systems are incomplete, but they do not suggest ways to measure the power of formal axiomatic systems, to rank their degree of completeness or incompleteness. Actually, Post's concept of a simple set contains the germ of the information-theoretic versions of Gödel's theorem that I will give later, but this is only visible in retrospect. One could somehow choose a particular simple r.e. set X and rank formal axiomatic systems according to how many different theorems of the form " N is not in X " are provable. Here are three other quantitative versions of Gödel's incompleteness theorem which do sort of fall within the scope of traditional methods.

Consider a particular formal axiomatic system in which it is possible to talk about total recursive functions (computable functions which have a natural number as value for each natural number input) and their running time computational complexity. It is possible to construct a total recursive

function which grows more quickly than any function which is provably total recursive in the formal axiomatic system. It is also possible to construct a total recursive function which takes longer to compute than any provably total recursive function. That is to say, a computer program which produces a natural number output and then halts whenever it is given a natural number input, but this cannot be proved in the formal axiomatic system, because the program takes too long to produce its output.

It is also fun to use constructive transfinite ordinal numbers (Hofstadter 1979) to measure the power of formal axiomatic systems. A constructive ordinal is one which can be obtained as the limit from below of a computable sequence of smaller constructive ordinals. One measures the power of a formal axiomatic system by the first constructive ordinal which cannot be proved to be a constructive ordinal within the system. This is like the paradox of the first unmentionable or undefinable ordinal number (Russell, 1967)!

Before turning to information-theoretic incompleteness theorems, I must first explain the basic concepts of algorithmic information theory (Chaitin, 1975b, 1977, 1982).

3. ALGORITHMIC INFORMATION THEORY

Algorithmic information theory focuses on individual objects rather than on the ensembles and probability distributions considered in Claude Shannon and Norbert Wiener's information theory. How many bits does it take to describe how to compute an individual object? In other words, what is the size in bits of the smallest program for calculating it? It is easy to see that since general-purpose computers (universal Turing machines) can simulate each other, the choice of computer as yardstick is not very important and really only corresponds to the choice of origin in a coordinate system.

The fundamental concepts of this new information theory are: algorithmic information content, joint information, relative information, mutual information, algorithmic randomness, and algorithmic independence. These are defined roughly as follows.

The algorithmic information content $I(X)$ of an individual object X is defined to be the size of the smallest program to calculate X . Programs must be self-delimiting so that subroutines can be combined by concatenating them. The joint information $I(X, Y)$ of two objects X and Y is defined to be the size of the smallest program to calculate X and Y simultaneously. The relative or conditional information content $I(X|Y)$ of X given Y is defined to be the size of the smallest program to calculate X from a minimal program for Y .

Note that the relative information content of an object is never greater than its absolute information content, for being given additional information can only help. Also, since subroutines can be concatenated, it follows that joint information is subadditive. That is to say, the joint information content is bounded from above by the sum of the individual information contents of the objects in question. The extent to which the joint information is less than this sum leads to the next fundamental concept, mutual information.

The mutual information content $I(X:Y)$ measures the commonality of X and Y : it is defined as the extent to which knowing X helps one to calculate Y , which is essentially the same as the extent to which knowing Y helps one to calculate X , which is also the same as the extent to which it is cheaper to calculate them together than separately. That is to say, $I(X:Y) = I(X) - I(X|Y) = I(Y) - I(Y|X) = I(X) + I(Y) - I(X, Y)$. Note that this implies that $I(X, Y) = I(X) + I(Y|X) = I(Y) + I(X|Y)$.

I can now define two very fundamental and philosophically significant notions: algorithmic randomness and algorithmic independence. These concepts are, I believe, quite close to the intuitive notions that go by the same name, namely, that an object is chaotic, typical, unnoteworthy, without structure, pattern, or distinguishing features, and is irreducible information, and that two objects have nothing in common and are unrelated.

Consider, for example, the set of all N -bit long bit strings. Most such strings S have $I(S)$ approximately equal to N plus $I(N)$, which is N plus the algorithmic information contained in the base-two numeral for N , which is equal to N plus order of $\log N$. No N -bit long S has information content greater than this. A few have less information content; these are strings with a regular structure or pattern. Those strings S of a given size having greatest information content are said to be random or patternless or algorithmically incompressible. The cutoff between random and nonrandom is somewhere around $I(S)$ equal to N if the string S is N bits long.

Similarly, an infinite binary sequence such as the base-two expansion of pi is random if and only if all its initial segments are random, that is, if and only if there is a constant C such that no initial segment has information content less than C bits below its length. Of course, pi is the extreme opposite of a random string: it takes only $I(N)$ which is order of $\log N$ bits to calculate pi's first N bits. But the probability that an infinite sequence obtained by independent tosses of a fair coin is algorithmically random is unity.

Two strings are algorithmically independent if their mutual information is essentially zero, more precisely, if their mutual information is as small as possible. Consider, for example, two arbitrary strings X and Y each N bits in size. Usually, X and Y will be random to each other, excepting the fact that they have the same length, so that $I(X:Y)$ is approximately equal

to $I(N)$. In other words, knowing one of them is no help in calculating the other, excepting that it tells one the other string's size.

To illustrate these ideas, let me give an information-theoretic proof that there are infinitely many prime numbers (Chaitin, 1979). Suppose on the contrary that there are only finitely many primes, in fact, K of them. Consider an algorithmically random natural number N . On the one hand, we know that $I(N)$ is equal to $\log_2 N + \text{order of } \log \log N$, since the base-two numeral for N is an algorithmically random $(\log_2 N)$ -bit string. On the other hand, N can be calculated from the exponents in its prime factorization, and vice versa. Thus $I(N)$ is equal to the joint information of the K exponents in its prime factorization. By subadditivity, this joint information is bounded from above by the sum of the information contents of the K individual exponents. Each exponent is of order $\log N$. The information content of each exponent is thus of order $\log \log N$. Hence $I(N)$ is simultaneously equal to $\log_2 N + O(\log \log N)$ and less than or equal to $K O(\log \log N)$, which is impossible.

The concepts of algorithmic information theory are made to order for obtaining quantitative incompleteness theorems, and I will now give a number of information-theoretic proofs of Gödel's theorem (Chaitin, 1974a, 1974b, 1975a, 1977, 1982; Chaitin and Schwartz, 1978; Gardner, 1979).

4. INFORMATION-THEORETIC PROOFS OF GÖDEL'S THEOREM

I propose that we consider a formal axiomatic system to be a computer program for listing the set of theorems, and measure its size in bits. In other words, the measure of the size of a formal axiomatic system that I will use is quite crude. It is merely the amount of space it takes to specify a proof-checking algorithm and how to apply it to all possible proofs, which is roughly the amount of space it takes to be very precise about the alphabet, vocabulary, grammar, axioms, and rules of inference. This is roughly proportional to the number of pages it takes to present the formal axiomatic system in a textbook.

Here is the first information-theoretic incompleteness theorem. Consider an N -bit formal axiomatic system. There is a program of size N which does not halt, but one cannot prove this within the formal axiomatic system. On the other hand, N bits of axioms can permit one to deduce precisely which programs of size less than N halt and which ones do not. Here are two different N -bit axioms which do this. If God tells one how many different programs of size less than N halt, this can be expressed as an N -bit base-two numeral, and from it one could eventually deduce which of these programs halt and which do not. An alternative divine revelation would be

knowing that program of size less than N which takes longest to halt. (In the current context, programs have all input contained within them.)

Another way to thwart an N -bit formal axiomatic system is to merely toss an unbiased coin slightly more than N times. It is almost certain that the resulting binary string will be algorithmically random, but it is not possible to prove this within the formal axiomatic system. If one believes the postulate of quantum mechanics that God plays dice with the universe (Albert Einstein did not), then physics provides a means to expose the limitations of formal axiomatic systems. In fact, within an N -bit formal axiomatic system it is not even possible to prove that a particular object has algorithmic information content greater than N , even though almost all (all but finitely many) objects have this property.

The proof of this closely resembles G. G. Berry's paradox of "the first natural number which cannot be named in less than a billion words," published by Russell at the turn of the century (Russell, 1967). The version of Berry's paradox that will do the trick is "that object having the shortest proof that its algorithmic information content is greater than a billion bits." More precisely, "that object having the shortest proof within the following formal axiomatic system that its algorithmic information content is greater than the information content of the formal axiomatic system: . . .," where the dots are to be filled in with a complete description of the formal axiomatic system in question.

By the way, the fact that in a given formal axiomatic system one can only prove that finitely many specific strings are random, is closely related to Post's notion of a simple r.e. set. Indeed, the set of nonrandom or compressible strings is a simple r.e. set. So Berry and Post had the germ of my incompleteness theorem!

In order to proceed, I must define a fascinating algorithmically random real number between zero and one, which I like to call omega (Chaitin, 1975b; Gardner, 1979). Omega is a suitable subject for worship by mystical cultists, for as Charles Bennett (Gardner, 1979) has argued persuasively, in a sense omega contains all constructive mathematical truth, and expresses it as concisely and compactly as possible. Knowing the numerical value of omega with N bits of precision, that is to say, knowing the first N bits of omega's base-two expansion, is another N -bit axiom that permits one to deduce precisely which programs of size less than N halt and which ones do not.

Omega is defined as the halting probability of whichever standard general-purpose computer has been chosen, if each bit of its program is produced by an independent toss of a fair coin. To Turing's theorem in recursive function theory that the halting problem is unsolvable, there corresponds in algorithmic information theory the theorem that the base-two expansion of omega is algorithmically random. Therefore it takes N bits of

axioms to be able to prove what the first N bits of ω are, and these bits seem completely accidental like the products of a random physical process. One can therefore measure the power of a formal axiomatic system by how much of the numerical value of ω it is possible to deduce from its axioms. This is sort of like measuring the power of a formal axiomatic system in terms of the size in bits of the shortest program whose halting problem is undecidable within the formal axiomatic system.

It is possible to dress this incompleteness theorem involving ω so that no direct mention is made of halting probabilities, in fact, in rather straightforward number-theoretic terms making no mention of computer programs at all. ω can be represented as the limit of a monotone increasing computable sequence of rational numbers. Its N th bit is therefore the limit as T tends to infinity of a computable function of N and T . Thus the N th bit of ω can be expressed in the form $\text{Exists } X \text{ Forall } Y$ (computable predicate of X , Y , and N). Complete chaos is only two quantifiers away from computability! ω can also be expressed via a polynomial P in, say, one hundred variables, with integer coefficients and exponents (Davis et al., 1976): the N th bit of ω is a 1 if and only if there are infinitely many natural numbers K such that the equation $P(N, K, X_1, \dots, X_{98}) = 0$ has a solution in natural numbers.

Of course, ω has the very serious problem that it takes much too long to deduce theorems from it, and this is also the case with the other two axioms we considered. So the ideal, perfect mathematical axiom is in fact useless! One does not really want the most compact axiom for deducing a given set of assertions. Just as there is a trade-off between program size and running time, there is a trade-off between the number of bits of axioms one assumes and the size of proofs. Of course, random or irreducible truths cannot be compressed into axioms shorter than themselves. If, however, a set of assertions is not algorithmically independent, then it takes fewer bits of axioms to deduce them all than the sum of the number of bits of axioms it takes to deduce them separately, and this is desirable as long as the proofs do not get too long. This suggests a pragmatic attitude toward mathematical truth, somewhat more like that of physicists.

Ours has indeed been a long stroll through a gallery of incompleteness theorems. What is the conclusion or moral? It is time to make a final statement about the meaning of Gödel's theorem.

5. THE MEANING OF GÖDEL'S THEOREM

Information theory suggests that the Gödel phenomenon is natural and widespread, not pathological and unusual. Strangely enough, it does this via counting arguments, and without exhibiting individual assertions which are

true but unprovable! Of course, it would help to have more proofs that particular interesting and natural true assertions are not demonstrable within fashionable formal axiomatic systems.

The real question is this: Is Gödel's theorem a mandate for revolution, anarchy, and license?! Can one give up after trying for two months to prove a theorem, and add it as a new axiom? This sounds ridiculous, but it is sort of what number theorists have done with Bernhard Riemann's zeta conjecture (Polya, 1959). Of course, two months is not enough. New axioms should be chosen with care, because of their usefulness and large amounts of evidence suggesting that they are correct, in the same careful manner, say, in practice in the physics community.

Gödel himself has espoused this view with remarkable vigor and clarity, in his discussion of whether Cantor's continuum hypothesis should be added to set theory as a new axiom (Gödel, 1964):

...even disregarding the intrinsic necessity of some new axiom, and even in case it has no intrinsic necessity at all, a probable decision about its truth is possible also in another way, namely, inductively by studying its "success." Success here means fruitfulness in consequences, in particular in "verifiable" consequences, i.e., consequences demonstrable without the new axiom, whose proofs with the help of the new axiom, however, are considerably simpler and easier to discover, and make it possible to contract into one proof many different proofs. The axioms for the system of real numbers, rejected by intuitionists, have in this sense been verified to some extent, owing to the fact that analytical number theory frequently allows one to prove number-theoretical theorems which, in a more cumbersome way, can subsequently be verified by elementary methods. A much higher degree of verification than that, however, is conceivable. There might exist axioms so abundant in their verifiable consequences, shedding so much light upon a whole field, and yielding such powerful methods for solving problems (and even solving them constructively, as far as that is possible) that, no matter whether or not they are intrinsically necessary, they would have to be accepted at least in the same sense as any well-established physical theory.

Later in the same discussion Gödel refers to these ideas again:

It was pointed out earlier... that, besides mathematical intuition, there exists another (though only probable) criterion of the truth of mathematical axioms, namely their fruitfulness in mathematics and, one may add, possibly also in physics... The

simplest case of an application of the criterion under discussion arises when some...axiom has number-theoretical consequences verifiable by computation up to any given integer.

Gödel also expresses himself in no uncertain terms in a discussion of Russell's mathematical logic (Gödel, 1964):

The analogy between mathematics and a natural science is enlarged upon by Russell also in another respect...axioms need not be evident in themselves, but rather their justification lies (exactly as in physics) in the fact that they make it possible for these "sense perceptions" to be deduced...I think that...this view has been largely justified by subsequent developments, and it is to be expected that it will be still more so in the future. It has turned out that solution of certain arithmetical problems requires the use of assumptions essentially transcending arithmetic...Furthermore it seems likely that for deciding certain questions of abstract set theory and even for certain related questions of the theory of real numbers new axioms based on some hitherto unknown idea will be necessary. Perhaps also the apparently unsurmountable difficulties which some other mathematical problems have been presenting for many years are due to the fact that the necessary axioms have not yet been found. Of course, under these circumstances mathematics may lose a good deal of its "absolute certainty;" but, under the influence of the modern criticism of the foundations, this has already happened to a large extent...

I end as I began, with a quotation from Weyl (1949): "A truly realistic mathematics should be conceived, in line with physics, as a branch of the theoretical construction of the one real world, and should adopt the same sober and cautious attitude toward hypothetic extensions of its foundations as is exhibited by physics."

6. DIRECTIONS FOR FUTURE RESEARCH

a. Prove that a famous mathematical conjecture is unsolvable in the usual formalizations of number theory. Problem: if Pierre Fermat's "last theorem" is undecidable then it is true, so this is hard to do.

b. Formalize all of college mathematics in a practical way. One wants to produce textbooks that can be run through a practical formal proof checker and that are not too much larger than the usual ones. LISP (Levin, 1974) and SETL (Dewar et al., 1981) might be good for this.

c. Is algorithmic information theory relevant to physics, in particular, to thermodynamics and statistical mechanics? Explore the thermodynamics of computation (Bennett, 1982) and determine the ultimate physical limitations of computers.

d. Is there a physical phenomenon that computes something noncomputable? Contrariwise, does Turing's thesis that anything computable can be computed by a Turing machine constrain the physical universe we are in?

e. Develop measures of self-organization and formal proofs that life must evolve (Chaitin, 1979; Eigen and Winkler, 1981; von Neumann, 1966).

f. Develop formal definitions of intelligence and measures of its various components; apply information theory and complexity theory to AI.

REFERENCES

Let me give a few pointers to the literature. The following are my previous publications on Gödel's theorem: Chaitin, 1974a, 1974b, 1975a, 1977, 1982; Chaitin and Schwartz, 1978. Related publications by other authors include Davis, 1978; Gardner, 1979; Hofstadter, 1979; Levin, 1974; Post, 1965. For discussions of the epistemology of mathematics and science, see Einstein, 1944, 1954; Feynman, 1965; Gödel, 1964; Polya, 1959; von Neumann, 1956, 1963; Taub, 1961; Weyl, 1946, 1949.

- Bell, E. T. (1951). *Mathematics, Queen and Servant of Science*, McGraw-Hill, New York.
- Bennett, C. H. (1982). The thermodynamics of computation—a review, *International Journal of Theoretical Physics*, **21**, 905–940.
- Chaitin, G. J. (1974a). Information-theoretic computational complexity, *IEEE Transactions on Information Theory*, **IT-20**, 10–15.
- Chaitin, G. J. (1974b). Information-theoretic limitations of formal systems, *Journal of the ACM*, **21**, 403–424.
- Chaitin, G. J. (1975a). Randomness and mathematical proof, *Scientific American*, **232** (5) (May 1975), 47–52. (Also published in the French, Japanese, and Italian editions of *Scientific American*.)
- Chaitin, G. J. (1975b). A theory of program size formally identical to information theory, *Journal of the ACM*, **22**, 329–340.
- Chaitin, G. J. (1977). Algorithmic information theory, *IBM Journal of Research and Development*, **21**, 350–359, 496.
- Chaitin, G. J., and Schwartz, J. T. (1978). A note on Monte Carlo primality tests and algorithmic information theory, *Communications on Pure and Applied Mathematics*, **31**, 521–527.
- Chaitin, G. J. (1979). Toward a mathematical definition of 'life', in *The Maximum Entropy Formalism*, R. D. Levine and M. Tribus (eds.), MIT Press, Cambridge, Massachusetts, pp. 477–498.
- Chaitin, G. J. (1982). Algorithmic information theory, *Encyclopedia of Statistical Sciences*, Vol. 1, Wiley, New York, pp. 38–41.
- Cole, C. A., Wolfram, S., et al. (1981). *SMP: a symbolic manipulation program*, California Institute of Technology, Pasadena, California.
- Courant, R., and Robbins, H. (1941). *What is Mathematics?*, Oxford University Press, London.

- Davis, M., Matijasevic, Y., and Robinson, J. (1976). Hilbert's tenth problem. Diophantine equations: positive aspects of a negative solution, in *Mathematical Developments Arising from Hilbert Problems, Proceedings of Symposia in Pure Mathematics*, Vol. XXVII, American Mathematical Society, Providence, Rhode Island, pp. 323–378.
- Davis, M. (1978). What is a computation?, in *Mathematics Today: Twelve Informal Essays*, L. A. Steen (ed.), Springer-Verlag, New York, pp. 241–267.
- Dewar, R. B. K., Schonberg, E., and Schwartz, J. T. (1981). *Higher Level Programming: Introduction to the Use of the Set-Theoretic Programming Language SETL*, Courant Institute of Mathematical Sciences, New York University, New York.
- Eigen, M., and Winkler, R. (1981). *Laws of the Game*, Knopf, New York.
- Einstein, A. (1944). Remarks on Bertrand Russell's theory of knowledge, in *The Philosophy of Bertrand Russell*, P. A. Schilpp (ed.), Northwestern University, Evanston, Illinois, pp. 277–291.
- Einstein, A. (1954). *Ideas and Opinions*, Crown, New York, pp. 18–24.
- Feynman, R. (1965). *The Character of Physical Law*, MIT Press, Cambridge, Massachusetts.
- Gardner, M. (1979). The random number omega bids fair to hold the mysteries of the universe, Mathematical Games Dept., *Scientific American*, **241** (5) (November 1979), 20–34.
- Gödel, K. (1964). Russell's mathematical logic, and What is Cantor's continuum problem?, in *Philosophy of Mathematics*, P. Benacerraf and H. Putnam (eds.), Prentice-Hall, Englewood Cliffs, New Jersey, pp. 211–232, 258–273.
- Hofstadter, D. R. (1979). *Gödel, Escher, Bach: an Eternal Golden Braid*, Basic Books, New York.
- Levin, M. (1974). *Mathematical Logic for Computer Scientists*, MIT Project MAC report MAC TR-131, Cambridge, Massachusetts.
- Polya, G. (1959). Heuristic reasoning in the theory of numbers, *American Mathematical Monthly*, **66**, 375–384.
- Post, E. (1965). Recursively enumerable sets of positive integers and their decision problems, in *The Undecidable: Basic Papers on Undecidable Propositions, Unsolvability Problems and Computable Functions*, M. Davis (ed.), Raven Press, Hewlett, New York, pp. 305–337.
- Russell, B. (1967). Mathematical logic as based on the theory of types, in *From Frege to Gödel: A Source Book in Mathematical Logic, 1879–1931*, J. van Heijenoort (ed.), Harvard University Press, Cambridge, Massachusetts, pp. 150–182.
- Taub, A. H. (ed.) (1961). *J. von Neumann — Collected Works*, Vol. I, Pergamon Press, New York, pp. 1–9.
- von Neumann, J. (1956). The mathematician, in *The World of Mathematics*, Vol. 4, J. R. Newman (ed.), Simon and Schuster, New York, pp. 2053–2063.
- von Neumann, J. (1963). The role of mathematics in the sciences and in society, and Method in the physical sciences, in *J. von Neumann — Collected Works*, Vol. VI, A. H. Taub (ed), McMillan, New York, pp. 477–498.
- von Neumann, J. (1966). *Theory of Self-Reproducing Automata*, A. W. Burks (ed.), University of Illinois Press, Urbana, Illinois.
- Weyl, H. (1946). Mathematics and logic, *American Mathematical Monthly*, **53**, 1–13.
- Weyl, H. (1949). *Philosophy of Mathematics and Natural Science*, Princeton University Press, Princeton, New Jersey.
- Wilf, H. S. (1982). The disk with the college education, *American Mathematical Monthly*, **89**, 4–8.