SANTA FE INSTITUTE

On Limits

John L. Casti and Joseph F. Traub

94-10-056

This research was supported in part by grants to the Santa Fe Institute, including core funding from the John D. and Catherine T. MacArthur Foundation; the National Science Foundation (PHY 9021437); and the U.S. Department of Energy (DE-FG05-88ER25054).

About the Santa Fe Institute

The Santa Fe Institute (SFI) is a multidisciplinary graduate research and teaching institution formed to nurture research on complex systems and their simpler elements. A private, independent institution, SFI was founded in 1984. Its primary concern is to focus the tools of traditional scientific disciplines and emerging new computer resources on the problems and opportunities that are involved in the multidisciplinary study of complex systems—those fundamental processes that shape almost every aspect of human life. Understanding complex systems is critical to realizing the full potential of science, and may be expected to yield enormous intellectual and practical benefits.

All Santa Fe Institute publications carry the SFI logo on the front cover, which is based on a Mimbres pottery design (circa A.D. 950-1150), drawn by Betsy Jones.

ON LIMITS

Edited by

John L. Casti
Santa Fe Institute
Santa Fe, New Mexico, USA

AND

Joseph F. Traub
Columbia University
New York, NY, USA

INTRODUCTION

At the *Limits to Scientific Knowledge* workshop held at the Santa Fe Institute during the period May 24–26, 1994, the participants agreed that it would be useful to set up an on-going teleconference to keep in touch on the many matters raised during the Workshop. As the first stage of this teleconference, participants agreed to submit one-page statements of their position(s) on aspects of limits to knowledge. This document contains these submissions in their raw, "unexpurgated" and unedited form. These initial statements sparked off additional discussion and debate on the general topic of the Workshop. The Appendix to this document contains these subsequent postings to the teleconference as of July 10, 1994.

Contents

26–May–94

Flow Past a Cylinder: An Example of Limited Scientific Knowledge

LEE A. SEGEL

I like to approach science by means of selected particular examples. Consider then the (two-dimensional) flow of a viscous incompressible fluid past a circular cylinder, or equivalently "flow past a circle." If we imagine the fluid to be unbounded, the problem is characterized by a single dimensionless parameter, the Reynolds number Re, which is the fluid speed U at infinity times the cylinder radius R divided by the fluid viscosity nu. Imagine a sequence of problems, at successively higher fixed values of Re. For Re somewhat greater than 10, a thin *boundary layer* forms near the cylinder, where the tangential velocity rapidly varies from a value of magnitude U to the value zero at the cylinder. The *wake* behind the cylinder has a structure similar to the boundary layer. At higher values of Re *instability* sets in, the symmetry is broken and an array of vortices (*vortex street*) forms aft of the cylinder. For higher Re, the overall gemoetry becomes very unsteady and the details of the flow appear *random*, and is termed *turbulent*.

The various italicized words are concepts with the aid of which the broad outlines of flow development can be understood. These concepts emerged from comparing experiment with analytic and numerical solutions to a well-accepted mathematical model for viscous flow, the Navier-Stokes (NS) equations. We will limit our further remarks to the behavior of relevant solutions to these equations, although it conceivable that in the future better equations will be found that will be free of the difficulties attached to solving the NS equations.

If more detailed aspects of the flow are sought, at present there is no recourse but to resort to numerical solutions of the NS equations. Such solutions yield results that appear in broad agreement with observation for Reynolds numbers of magnitude $Re = 1000$ but not much larger. Even in these instances the famous "exponential divergence" phenomenon of chaotic flows makes it prohibitive to calculate the paths of individual fluid particles for more than a rather short time. Nonetheless "global" or "averaged" features of the flow, such as the force exerted on the cylinder, can be calculated.

Here are two sample questions whose answers would add to our understanding of knowledge limitations in this example. (i) For $Re = 1000$, how would computer power have to be multiplied so that we could calculate a particle path with 1times. (ii) How would computer power have to be multiplied if we wished to calculate the force on the cylinder for $Re = 10^n$ for $n = 6$ and $n = 9$ (which are "reasonable" values for certain applications).

More generally, the fact that climate seems harder to predict than weather indicates that functionals of solutions may not necessarily be easier to calculate than solution details (although for the longer time scales required for climate forecasts, new effects may complicate the situation compared to weather forecasts). Can various examples be found of functionals that are and are not easier to calculate than details?

A final point. If we wished to prove that certain matters are unknowable in a practical sense for the NS equations at large values of the Reynolds number, then we would have to reckon with the fact that analytic methods are particularly powerful precisely in this parameter domain. To my knowledge, there are no analytical results that provide an asymptotic formula for, say, the drag on the cylinder as $Re \to \infty$. It would be nice to find such a formula. Or to prove that no such formula exists.

02–Jun–94

Impossibility Results in Mathematics

G. J. Chaitin

Call a program "elegant" if no smaller program produces the same output. You can't prove that a program is elegant. More precisely, N bits of axioms are needed to prove that an N-bit program is elegant.

Consider the binary representation of halting probability Ω. The number Ω is the probability that a program chosen at random halts. You can't prove what the bits of Ω are. More precisely, N bits of axioms are needed to determine N bits of Ω.

I have constructed a monster algebraic equation

$$P(K, X, Y, Z, \ldots) = 0.$$

Vary the parameter K and ask whether this equation has finitely or infinitely many whole-number solutions. N bits of axioms are needed to be able to settle N cases.

These striking examples show that sometimes one must put more into a set of axioms in order to get more out. They are accidental mathematical assertions that are true for no reason, that is, irreducible. The only way to prove them is essentially to add them as new axioms. Thus in this extreme case one gets out of a set of axioms only what one puts in.

How are these results proved? The basic idea is contained in the Berry paradox of "the first positive integer that cannot be specified in less than a billion words." This leads one to define the program-size complexity of something to be the size in bits of the smallest program that calculates it. This is the basic concept of algorithmic information theory (*AIT*). *AIT* is an elegant theory of complexity, perhaps the most developed of all such theories, but one should recall von Neumann's remark that pure mathematics is easy compared to the real world. *AIT* provides the correct complexity concept for metamathematics, but it is not the correct complexity concept for physics, biology, or economics.

Program-size complexity in *AIT* is analogous to entropy in statistical mechanics. Just as thermodynamics gives limits on heat engines, *AIT* gives limits on formal axiomatic systems.

I have recently reformulated *AIT*. Up to now, the best version of *AIT* studied the size of programs in a computer programming language that was not actually usable. Now I obtain the correct program-size complexity measure from a powerful and easy to use programming language. This language is a version of *LISP*, and I have written an interpreter for it in *C*. A summary of this new work is available as IBM Research Report RC–19553 "The limits of mathematics," which I am expanding into a book.

Bibliography

[1] G. J. Chaitin, "Randomness and complexity in pure mathematics," *International Journal of Bifurcation and Chaos* **4** (1994), pp. 3–15.

[2] G. J. Chaitin, "Responses to 'Theoretical Mathematics . . . ,'" *Bulletin of the American Mathematical Society* **30** (1994), pp. 181–182.

Some Questions for Future Research

(1) What is an interesting or natural mathematical question?

(2) How often is such a question independent of the usual axioms?

(3) Show that a classical open question such as the Riemann hypothesis is independent of the usual axioms.

(4) Should we take incompleteness seriously or is it a red herring?

(5) Is mathematics quasi-empirical? In other words, should mathematics be done more like physics is done?

Study of Intrinsic Limits to Knowledge

PIET HUT

Nearly four centuries ago, the need was recognized for a demarcation of the scientific method from ideological dogma, whether religious or political. Independent of one's personal convictions, a scientist was supposed to be led by experimental evidence alone in the construction of scientific theories. During the subsequent three centuries, this approach paid off: mathematics and the various natural sciences underwent a spectacular growth in depth and scope. Once liberated from external shackles (and funded generously by outside sources), it seemed that the only limitations on our knowledge in these areas were external, given by limits of time, energy and creativity.

There was, however, a type of unquestioned dogma that had been taken over unwittingly from the preceding scholastic period: the notion of a 'God's eye view', of an ideal limit point to which increasingly detailed knowledge could strive. And indeed, until early this century, nothing seemed to stand in the way of an ever-more-detailed understanding of what was objectively 'out there', waiting to be charted and analyzed. We now know that this notion has proved to be naive. After rejecting any external source of authority in our study of nature, we have found that science cannot provide its own version of an external view of nature, not even in principle.

The discovery of the quantum mechanical nature of the physical world showed that there are intrinsic limits to the amount of knowledge that can be extracted from measurements of a physical system. The subsequent discoveries of intrinsic limits to mathematical logic, starting with Gödel's incompleteness proofs, showed that even in the ideal world of mathematics we cannot assume a 'God's eye view' of the whole terrain as an ideal to strive towards.

All this came as a great surprise. In the last sixty years, we have learned to live with these types of limits in their respective domains, but most of the subsequent work has been area-specific. For example, in astrophysics we are faced with the fact that only part of the universe is contained within our past light cone, making the rest of the universe (as yet) inaccessible. In addition, any cosmological theory can be tested only on the one universe we happen to live in, thereby providing an intrinsic constraint on experimental verification of probability distributions in such a theory.

During the workshop on *Limits to Scientific Knowledge*, we have heard many other examples, ranging from limits to practical computability of complex problems in mathematics to limits to the universal concepts our minds have learned to work with during the evolutionary adaptation to our surroundings. A prime example of the latter is the question about the hierarchical structure of most of our knowledge: to what extent is this a reflection of how the world is built up, and to what extent is this a result of our limitations in recognizing non-hierarchical forms of organization?

A widely interdisciplinary attempt to link up some of these diverse investigations would seem timely, judging from the excitement generated in the many discussions during the workshop. Just as the area of 'cognitive science' sprung up naturally across the interfaces between a number of different fields, we may witness the formation of a new area, 'limits to cognition', or 'knowledge of (limits to) knowledge'. And to stress the intrinsic nature of these limits, a name for a pilot program could be the 'Study of Intrinsic Limits to Knowledge'. At least its acronym, SILK, would be pronounceable.

There are many directions that such a program could take. For example, one of our meetings could take the form of a summer school, perhaps a two-week program in which a limited number of topics are discussed more in-depth, in the form of a series of lectures. The presence of graduate students could add significantly in generating new and fresh ideas for further explorations.

Limits in Economics

ALFRED P. NORMAN

One concept which the participants of the conference were in agreement was that models of the physical world should be computable, at least as an approximation. In the social sciences this requirement should be strengthened to assert that models of a social agent should be computable by the computational resources of that social agent. The imposition of this requirement of agent computability will result in a paradigm shift in economics.

Almost all microeconomic models are based on what Simon calls substantive rationality. What this means is that economists examine the consequences of optimizing behavior without asking the question how the economic agent performs the optimization. The reason for this approach is that the first order conditions give the consumer demand functions for utility maximization and factor demand functions for profit maximization.

The requirement that an economic agent model be agent computable should be applied to the procedures used by economic agents, what Simon calls procedure rationality. In creating such models it is desirable that such models reflect the economic concerns of the 1990s not the 1870s when substantive rational models were originally developed. In the 1870s it was assumed that the consumer knew his preferences and the producer knew his production function; consequently the concern was resource allocation with fixed technology and known preferences. With a high rate of technological change, this is not a valid assumption today. Currently, worldwide economic competition is based on the discovery of new knowledge, the invention of new products from that knowledge, and a constant innovation of more efficient production processes. Procedural models of economic agents should reflect these concerns.

The need to learn to improve the production process can greatly increase its computational complexity. For example, the introduction of an unknown parameter into the production function of a simple profit maximization problem can increase its computation complexity from 0 to ∞. Moreover, it is not obvious that good agent computable approximations exist in the unknown parameter case. Consequently, procedural models must anticipate that agents can make mistakes, such as GMs botched attempts to automate its production processes. Mistakes are impossible under substantive rationality by definition.

Likewise, procedural models of consumers are needed to create good information policy. Under perfect competition, the substantive consumer is assumed to possess all information, which he or she can process instantaneously at zero resource cost. Under a restriction of agent computability, a consumer would not process all information even if he or she had it. In addition, in metropolitan markets individuals must decide among a potentially very large number of alternatives. Psychologist have categorized many decision rules used by consumers. Some of these decision rules are sublinear in their computational complexity; however they are prone to error. As markets move to computer networks, a good model of the procedural consumer is needed to define informational policy for decision makers. There is a fundamental conflict between proprietary rights, privacy and efficiency.

References

1. Norman, A., 1993 *Informational Society, An economic theory of discovery, invention and innovation*, (Kluwer: Boston).

2. Norman, A., 1994, "Risk, Uncertainty, and Complexity," *Journal of Economic Dynamics and Control*, 18, 231–249.

3. Norman, A., 1994, "Technology, Complexity and Consumer Choice," Manuscript in fight status with referees.

4. Norman, A., 1994, "Computability, Complexity and Economics," *Computational Economics*, 7, 1–21.

06–Jun–94

Limits to Scientific Knowledge

ROGER SHEPARD

1. The history of science reveals a pattern (most fully, in the most highly developed science—physics) in which wider and wider ranges of phenomena have been subsumed under smaller and smaller sets of mutually consistent principles. (Even when we deal with highly complex systems, whether meteorological, evolutionary, neuronal, or economic, we secure scientific knowledge only to the extent that we discover general principles that, acting together, predict or explain aspects of the complex behavior we observe.) Scientific knowledge would therefore seem ultimately to be limited by any constraints on our ability to continue such a process toward a mutually consistent (probably hierarchical) system of universal principles.
2. The prevailing empiricistic account of science would be tenable if we had two or more theories that equally well accounted for everything we already know to be true. In that case, it might be only through the collection of more data that we could eliminate one or more of the competing theories. In fact, however, that case has never arisen. Even now, our current best theories—quantum mechanics and general relativity—are not mutually consistent. The construction of one self-consistent theory that subsumes those two theories would not so much seem to require more data as conceptual insight comparable, perhaps, to those that led to relativity and quantum mechanics themselves. But those insights were based on thought experiments already extending far beyond the speeds, energies, and sizes of our familiar macroscopic domain. Those insights also entailed increasingly abstract mathematical formalisms. Hence, further progress in theoretical physics will be subject to the as yet unformalized limitations of brains that have evolved in a three-dimensional locally Euclidean world of relatively slow-moving macroscopic objects.
3. The difficulty of finding a self-consistent theory to explain all the facts that we already know to be true about the world raises Einstein's question of whether, in establishing the laws of physics, "God had any choice." Some have suggested not (with Leibnitz) that this is "the best of all possible worlds" but, rather, that this may be the only possible world. The alternative—that the laws governing the world are merely arbitrary—would place limits on scientific knowledge that are inexplicable.
4. If, however, the final theory is determined by mathematical consistency, we confront the old, still unresolved question of the ontological status of mathematical truths: Do they exist prior to our discovery of them, or are they merely inventions of the the human mind? In accepting the latter alternative, one would seem in effect to be assuming something unlikely to be widely embraced by physicists—that the laws of physics can ultimately be understood only in terms of the laws of the mind.
5. Finally, even if we had a self-consistent theory that subsumed our present best scientific theories, it still would not explain everything we already know. There are two starting points for knowledge: There is the objective world (usually presupposed in science), within which

certain orderly phenomena (whether spectral lines on photographic plates or behaviors of other persons) are explained by positing unobserved structures and processes (whether atomic or mental). There is also one's own subjective experience, within which certain orderly correlations lead one to posit external structures and processes (including the physical world and other persons). Perhaps the most fundamental limit to knowledge is this irreducible objective-subjective complementarity. After many centuries of trying, no one has convincingly identified any necessary connection between an objective world (or its particles, waves, or space-time curvatures) and the subjective qualia of, say, pains, odors, or colors (or why short wavelengths yield the subjective experience blue, long wavelenghts red, and not the other way around). Can such limits ever be formalized?

Limits to Scientific Knowledge

ROLF LANDAUER

Since 1967 my papers have emphasized that there is likely to be a limit to the maximal number of computational steps available, in principle, in the universe. Thus, the algorithms invoked by our ordinary continuum mathematics are not executable, and therefore not suitable as a basis for the laws of physics. Algorithms for real use need to be restricted to calculations which can, at least in principle, be carried out. This point may be related to, but is distinct from, the more common observation that the physical universe is not properly described by the real number continuum. My views on this subject have received little attention, and only informal oral debate. That may have several sources, but I suspect that a key one relates to the fact that scientists pay attention to work which facilitates their next paper. Limits which seem far from a visible impact do not qualify.

Where does the knowledge we are considering reside? If it is to be used in decision making (social sciences) it may need to reside in the mind. But within the physical sciences the old emphasis on the observer may be excess baggage. In a day of genetic manipulation, scanning tunneling microscopy, spectroscopy of single atoms, etc., it is not clear that we should limit ourselves to information which has been passed back up to a macroscopic level.

In mathematics, physics, computer science, molecular biology, etc. we face large areas of potential knowledge against which the limitations we have discussed represent interesting islands of inaccessibility. By contrast, in the study of social structure, including economics, we face complex systems which are not only beset by chaos, but also controlled through occasional large and irregular "noise" bursts. ("Noise" is an influence arising from events not contained in our description of the system's dynamics, e.g., the effect of a presidential assassination on the economy.) Stressing the unpredictability of an economy is not very useful. Can we learn more about the kinds of things which may be predictable, and can we do so in advance of the actual detailed techniques for such predictions?

Undecidability and Biology

ROBERT ROSEN

To me, the basic question in biology, to which all others are subsidiary or collateral, is the one put most succinctly by the physicist Erwin Schrödinger: What is Life?

Any question becomes unanswerable, if we do not permit ourselves a uni- verse large enough to deal with the question. $Ax = B$ is generally unsolvable in a universe of positive integers. Likewise, generic angles become untrisectable, cubes unduplicatable, etc., in a universe limited by rulers and compasses.

I claim that the Gödelian noncomputability results are a symptom, arising within Mathematics itself, that we are trying to solve problems in too limited a universe of discourse. The limits in question are imposed in mathematics by an excess of "rigor"; in science, by cognate limitations of "objectivity" and "context independence." In both cases, our universes are limited, not by the demands of problems that need to be solved, but by extraneous standards of rigor. The result, in both cases, is a mindset of Reductionism, of only looking downwards towards subsystems, and never upwards and outwards.

In science, for instance, it seems patently obvious to us that, whatever living organisms are, they are material systems, special cases drawn from a larger, more generic class of nonliving inorganic ones. The game is thus to reduce, to express their novel properties in terms of those of inorganic subsystems merely subject to a list of additional conditions and restrictions. Indeed, one manifestation of this claim to the "objectivity" of Reduction is that one must never, ever, claim to learn anything new about matter from a study of organisms. This is but one of the many forms of the protean "Central Dogma," expressed here as a limitation on material nature itself.

I have long noted that, despite the profound differences between those material systems which are alive and those which are not, these differences have never been expressible in the form of a list, an explicit set of conditions which formally demarcate those material systems which are organisms from those which are not. Without such a list, Schrödinger's question, and biology itself, become unanswerable at best, meaningless at worst. So we must probe more deeply into what the quest for such a list actually connotes.

No such list means there is no algorithm, no decision procedure, whereby we can find organisms in a presumably larger universe of inorganic systems. It has of course never been demonstrated that there is no such list. But no one has ever found one, I take seriously the possibility that there is no list, no algorithm, no decision procedure, which finds us the organisms in a presumptively larger universe of. inorganic systems. This possibility is already a kind of non-computability assertion.

Indeed, the absence of lists or algorithms is a generally recurring theme in science and mathematics, one which reveals the nongenericity of the world of algorithms itself, a world too unstable (in a technical sense) to solve the real prob- lems. This was the upshot of the Gödel results from the very beginning.

It helps to recall the mathematical situation which Gödel inherited. It was a world still reeling from the discovery of "Non-Euclidean" geometries almost a century earlier, geometries without number which were just as consistent as Euclid was. It was a world reeling from paradoxes within Cantorian set theory. There had to be something to blame for all of this, something to be expunged, to make everything right again; something not rigorous enough, which had to be found and eradicated.

Bertrand Russell, among others, argued that the fault lay in "impredicative" definitions; vicious circles, and developed an elaborate and murky "theory of types" to replace them with predicative but equivalent counterparts. This was taken yet further by Hilbert and his school of "formalists," who argued that rigor lay entirely in syntax, and that the difficulties at the foundations of mathematics arose entirely from unextruded, semantic residues of meaning. For them, a mathematical term (e.g., "triangle") was not to be allowed any vestige of meaning, rather, there were to be formal production rules for manipulating "triangle" from one proposition to another. This drastic extrusion of semantics constituted true rigor; mathematics itself would be suspect as long as there was any vestige of meaning or semantics left in it. Hilbert sought this kind of "formalization" of all of mathematics, the "reduction" of mathematics to algorithms or lists.

It was this program which Gödel's results killed. Let us briefly recount what these results actually mean. They mean that a constructive universe, finitely generated, consisting of pure syntax, is too poor to do "mathematics" in. They mean that semantics and impredicativities and meanings are essential to mathematics, they cannot be replaced by more syntactic rules and more lists or algorithms. They mean that "mathematical systems" are generically unformalizable; hence it is the formalizable ones which are the rare special cases, and not the other way around. They mean that identifying "rigor" with formalizability makes most of mathematics unreachable.

I argue that biology teaches us the same is true about the material world. Roughly, that contemporary physics is to biology as number theory is to a formalization of it. Rather than an organism being just a standard material system plus a list of special conditions, an organism is rather a repository of meanings and impredicativities; more generic than an inorganic system rather than less. If this is so, then the Schrödinger question, and indeed biology itself, is not exclusively, or even mainly, an empirical science; empirics is to it as accounting is to number theory.

If this is so, then organisms possess noncomputable, unformalizable models. Such systems are what I call complex. The world of these systems is much larger and more generic than the simple world we inherit from Reductionism.

The main lesson from all this is that computability, in any sense, is not itself a Law of either Nature or of Mathematics. The noncomputability results, of which Goedel's was perhaps the first and most celebrated, are indicative of the troubles which arise when we try to make it such.

Limitology

OTTO E. RÖSSLER

The notion of "Maya's veil" in India, the phrase "the phenomena are a view of the hidden" by Anaxagoras, and the "flame sword" of the Bible are perhaps the oldest examples of limit statements. In modern science, Boscovich (1755) formulated the interface principle: "the impressions generated" are invariant under certain objective changes of the world (as when the earth is "breathing" along with the observer and all forces). Next, Maxwell (1872) predicted that from the inside of the world, the second law of thermodynamics cannot be violated by an attempt at exact observation (followed by sampling) of cold uranium gas atoms—to use a modern example—and that, therefore, micro motions are in general "impalpable" (p. 154 of his *Theory of Heat).* He thereby predicted quantum mechanics. Third, Bohr (1927) claimed that his "complementarity" defines the limits to what a classical internal observer of the world can hope to measure. Gödel's inaccessibility limit (inaccessibility in finitely many steps of certain implications of formal systems) came in the wake of Bohr's stronger (distortion-type) but less formalized limit. More recent limits are Moore's (1956) analogue to uncertainty in dissipative automata, Mackay's (1957) irrefutability of certain false statements (assertion of free will) in deterministic automata, Rosen's (1958) limit to self- reproduction in category theory, Lorenz's (1964) inacessibility principle in dynamical systems (butterfly effect), and—in the 1970s—finite combinatorial inaccessibility (NP-completeness) .

Do the two types of limit which are known so far (inaccessibility; distortion) justify the introduction of a new branch of science? Two points speak in favor of the idea. First, new limits become definable within the fold of the new paradigm. For example, ultraperspective which is inaccessible to animals plays a crucial role in both mathematical economics and science. The very notion of limits implies adopting two alternative perspectives simultaneously. Second, distortion-type limits lead to a "splitting" of the directly inaccessible (exo) world into many different internally valid (endo) worlds. Part of the objective features of the world must belong to the observer-specific interface ("subjective objectivity"). Thus observer-relative objective features of the world exist not only in relativity (with its observer-frame-specific mass and simultaneity), but also in quantum mechanics. An unexpected unification suggests itself ("micro relativity").

The main message of distortion-type limits is that there is only one way to regain an objective picture that does not depend on the observer: by making the observer explicit in an artificial-world approach. The computer acquires a new fundamental role.

I thank John Casti and Joe Traub for stimulation. Discussions with Oliver Morton on Maxwell, Rolf Landauer on a new counterfactual transluminality, Patrick Suppes on Kant and Chico Doria on consistency were helpful.

Final Comments on the Workshop "Limits to Scientific Knowledge"

E. ATLEE JACKSON

To have a clear discussion about these "limitations," it is necessary to have some working definition of "(natural) scientific knowledge." Whatever this definition may be, it will establish one class of limitations—the boundary between our experiences in the world and the discovery of sets of "correlated scientific observables," which can be replicated (or compared), and recorded for future reference. Questions related to the generalization of this boundary address this class of "limitations" (e.g., Hut, Rössler, Rosen, and Shepard; also parts of economics, psychology, and social areas).

A second class of limitations, which might be (!) of interest to "traditional" scientists, are related to those that arise within the restricted operations they use to obtain "scientific knowledge." Presently their sources of information are from physical observations (PO), mathematical models (MM), and computer experiments (CE), and these types of information are mutually distinct. Thus one subset of limitations arise in relating these disparate forms of information (encoding/decoding associations). This encoding/decoding process is necessary in any cyclic connection of PO with MM and/or CE, used in some form of "scientific method," which is intended to validate our understanding of particular physical phenomena. The PO are obtained from Mother Nature with the aid of some instruments, which maps a phenomenon onto a record.

Similarly, the MM Nature, involving some equations, requires some "analysis instruments" to yield a deduction (MM-observable). The genius of Poincare was to supply many of these new methods of analyses, when the historic analytical approachs failed. Also, if a CE is run (as in A-Life, Arthur's, and Epstein's studies), the resulting CE Nature, displayed on the computer screen or in data, requires new methods of analysis in order to produce "CE observables" (some objective, comprehensible record—without such CE observables, the bottom-up calculations reflect another form of microreductionism). It is these sets of observables that are encoded/decoded in any scientific method. The message here is that there are limitations implicit in the existing sets of "instruments" available for observing each of these "Natures." These limitations are real, but variable with the evolution of science, representing stimulating challenges to scientists.

Certainly a basic limitation that we have learned from chaos is that scientific knowledge is inhomogeneous in character, both as it relates to temporal extent, and to the particular observed phenomenon. Here we must pay attention to the finite character of observational information, both in PO and CE. The finer the scale of the phenomenon, the shorter will be the predictive period for the same informational input. A totally chaotic phenomenon (e.g., a Bernoulli shift) simply transcribes initial "information-in" proportionally into predicted "information-out." Thus certain average forms of information may be predicted much easier than detailed microscopic features (as noted by Segel, and used by Hut). This is particularly important to recognize in CEs, since they always involve round-off errors in their details (the "shadowing problem"), which may not, however, imply errors in averaged observables—indeed, they may more accurately represent the universal couplings between all systems at this level. Presently the nature of these limitations are not well understood.

A larger issue concerns establishing a program of science, to replace the nonscientific microreductionistic philosophy of this century, which will give some basis for unifying our understanding of phenomena at all levels of complexity. "Understanding" a phenomenon has always been accomplished in physics only by deductively relating it to interactions between rather simple "collective variables" (a "local" reduction-and-synthesis). This certainly suggests that our understanding of more complex systems will be limited to some form of hierarchical relationships—that the concepts of "fundamental laws," or "final theories," while possibly philosophically comforting or religiously engendered, have no scientific basis. As D. Boorstin noted: "The great obstacle of man is the illusion of knowledge ... The great significance of the voyages of discovery was their discovery of ignorance."

Knowledge, Randomness and Simulations

P.-M. BINDER

Every random number generator will fail in at least one application.
–D. Knuth

We do not know for sure whether the gods are ultimately watchmakers or dice-players. Consequently we model natural processes as either deterministic or random, the latter both for our mathematical convenience and in order to disguise our partial ignorance.

The second of the approaches above has been nagged constantly by foundational issues [1]; I will discuss one of them: the effect of artificially generated randomness in computer simulations. The latter play a central role in the development of the statistical micro-to-macro connection in physics: many microscopic models require simulations in which local variables (e.g., spins) are susceptible to change by thermal randomness, which is generated in the computer with simple algebraic formulas [2] that are iterated billions and billions of times.

(Statistical physics is a mature, established field. By comparison, little is known about similar but deterministic systems such as cellular automata which often exhibit persistent structures, universal computation and bizarre conserved quantities which cannot be satisfactorily studied with statistical methods).

Knuth [3] has reviewed a number of schemes for generating random sequences of numbers and tests of how random these sequences are. An important point is that these tests only look for specific types of correlations, and unexpected things can happen. For example, several random number generators were recently used in simulations of the two-dimensional Ising model, and compared with exact analytical solutions [4]. One of the generators (the widely used R250) yielded systematically and significantly incorrect results for the energy and specific heat. The errors were traced back to subtle correlations in the high-order bits of the random numbers, which were subsequently easily corrected.

In general we cannot expect to be so lucky. Exact solutions are only known for very few problems, and for the rest we are in the hands of the convenient but undependable procedures just outlined for simulations. In a few cases it is possible to reformulate the models in terms of probability transition matrices in such a way that the simulations become "exact" again [5].

But in general the scientific knowledge resulting from our Monte Carlo simulations is limited by the poorly understood nature of artificially generated randomness and by the tight interplay between the random sequences and model dynamics.

Typically the simulations offer evidence for (unknown or poorly understood) correlations induced by the microscopic model, and it is sobering that we cannot tell for sure which correlations come from the model and which come from the generator. For these purposes the results of standard tests for randomness are insufficient, and there is no systematic way to devise tests for particular applications. In such cases the best strategy is to explore the robustness of the results to changes in the parameters describing both the system and the random number generator, a procedure which presently has the status of an art form.

References

[1] L. Sklar, *Physics and Chance*. Cambridge, U.K.: Cambridge University Press, 1993.

[2] Note that the resulting sequences are algorithmically simple in the Chaitin-Kolmogorov sense.

[3]. D. Knuth, *The Art of Computer Programming*. Ch. 3 – Random Numbers. Reading, MA: Addison-Wesley, 1981. There is an especially interesting section about tests for randomness in short sequences.

[4] A. M. Ferrenberg, D. P. Landau, and Y. J. Wong, *Phys. Rev. Lett.*, 69, 3382 (1992).

[5] See S. Havlin and D. Ben-Avraham, *Adv. Phys.*, 36, 695 (1987) for a review.

13–Jun–94

On Undecidability and Incompleteness

F. A DORIA

Nah ist,
und schwer zu fassen der Gott.
Wo aber Gefahr ist,
waechst das Rettende auch.
(Hoelderlin, *Patmos*)
(I quote from memory, so mistakes are allowed.)

As far as I can tell both from my own impressions after perusing some of the recently published late texts by Gödel, and from the opinions of people who were in direct contact with him, Gödel wasn't so sure that Church's Thesis was the last word on computability. I know that there is an article by M. Davis which asserts the opposite, but his arguments never convinced me. It seems that Gödel looked upon his incompleteness theorem as some kind of limitation imposed by the current views on mathematics and its foundations.

Usually people try to circumvent the incompleteness phenomenon by adding new, stronger axioms to the already incomplete axiomatic theory. Incompleteness remains, but we are able to decide out of the new axioms several previously undecidable facts. However Gödel seemed to think that we might profitably *expand* the concept of provability and computability; this can be inferred from his 1958 proof of the consistency of arithmetic, where he clearly extends the idea of finitariness in order to attain his goal. Da Costa and myself have tentatively remarked that several 'mechanical' natural strict extensions of the current notion of computability are able to decide arithmetic sentences along the standard model. Our construction is, I guess, formally equivalent to Gödel's consistency proof, while growing out of a different intuition.

Yet there is another side to the decision problem. As I have been insisting whenever people equate incompleteness with some kind of formal limitation in mathematics, undecidable sentences can be coded as bifurcation points in vector fields. (See da Costa and Doria, "An undecidable Hopf bifurcation with an undecidable fixed point," to appear in the *Int. J. Theor. Phys.*) So, when we *decide* undecidable sentences because we are using stronger proof techniques, we are losing some of the degrees of freedom which are available to us in an incomplete theory. To summarize:

incompleteness = bifurcation points in theories.

different models for theories = different *representations* for the same syntactic object.

stronger proof techniques = 'interactions' (in the sense of physics) that restrict the degrees of freedom available in a theory.

The first point is already settled. The other two are still resilient, a pleasant speculation to me.

16–Jun–94

Why Undecidability is Too Important to be Left to Gödel

JOHN L. CASTI

How difficult is it to answer a question like "Where will the Dow Jones Industrial Average sit at the end of 1994?", or "What are the odds of the Dallas Cowboys repeating as the NFL champions?", or even"On what day will the Santa Fe Institute *really* move to its new campus off Hyde Park Road?" While none of these questions is particularly "scientific," at least in the usual sense of that term, they all have far more immediacy and importance for the daily lives of most of us than do far more scientifically respectable queries such as "What is minimal cost of touring the 50 state capitals?" or "What is the "complexity" of Champernowne's number?" The point here is that the traditional measures of computational and algorithmic complexity as used in mathematics and computer science don't even begin to characterize the difficulties inherent in trying to come up with believable, let alone correct, answers to these sorts of questions.

So there most certainly are limits to our knowledge—scientific or otherwise. But for the most part these limits are *not* the kind envisioned by Gödel, Turing, and their many able successors. Rather they are limits of a far subtler and more insidious kind: limits due to the ill-posed nature of most questions, limits stemming from the vagueness in natural human languages, limits coming from our inability to make error-free measurements, limits arising out of weaknesses in our standard modes of reasoning and, in general, limits generated by the inherently complicated mutual interactions between the many components of most natural and human systems.

So, in my view the most productive way to go about understanding whether or not any of these kinds of limits to our knowledge are permanent or merely temporary barriers to our understanding of the world we live in is to think about the development of new ways of looking at complexity in formal terms, ways that more faithfully mirror our informal ideas about what is and isn't "complex." These ways must necessarily include the crucially important fact that real-world complexity (whatever *that* might be) is not a context-free property of anything. Like truth, beauty, justice and "the good," complexity resides as much in the eye of the beholder as it in the beholden.

Whatever such a real-world measure of complexity turns out to involve, almost certainly will not be something as simple and naive as a single number characterizing the complexity of a given problem, question or situation. Rather, I envision it being some kind of multi-criterion, interactive type of measure, saying in effect that some kinds of problems are very complex indeed if you attack them by, say, deductive modes of reasoning, but evaporate away like trickle of water in the desert if you use other modes of reasoning like induction, abduction or (fill in your favorite "-ism" here). Sad to say, I have no such "grand unified theory of complexity" to offer up for your entertainment at the moment. But all my intuitions say that such a theory is not beyond hope. Furthermore, I feel strongly that only when a skeletal framework of this kind of theory emerges will we finally be on the way to banishing the ghost of Gödel from the scientific banquet table.

30–Jun–94

Upping the Ante

JOSEPH TRAUB

There are numerous impossibility results and conjectures in mathematics and computation. Major contributors include Gödel, Turing, Landauer, Chaitin, Matyasevich, Cook, and Karp.

Can we up the ante and prove that there are unanswerable questions in science that is, limits to scientific knowledge? A key word in that sentence is *prove*. Can we turn philosophical inquiry into proof?

Scientific questions do not come equipped with a mathematical model. Examples of scientific questions are:

- How did the matter left over from the Big Bang form stars, galaxies, and superclusters of galaxies?
- When will the universe stop expanding?
- How do we learn a language?
- What will be the average world temperature in 2001?

I noticed a tendency in our Workshop discussions and in these postings to invoke a particular mathematical model and to discuss undecidability or intractability for that model. An important issue is whether we can divorce unknowability of scientific questions from mathematical models. This suggests:

Issue 1. Can we consider scientific questions separate from mathematical models in our studies?

On the other hand, mathematical models are of interest in themselves. It seems to me that a mathematical model should satisfy the following twin desiderata:

- It should capture the essence of the scientific question (that's a scientific judgment).
- It should be computationally tractable.

This observation can be used in two ways. The first is that if we can prove that there exist scientific questions such that every mathematical formulation is intractable in every setting (that is, no matter how much we weaken the assurance) then we have a means for establishing the existence of unanswerable questions.

That's a big *If*. Although computer scientists routinely consider all possible algorithms (including those already known and also those that may exist in principle) scientists and mathematicians do not consider all possible mathematical models that capture the essence of a scientific question. This leads me to propose:

Issue 2. Can we fruitfully consider the notion of all mathematical models that capture the essence of a scientific question?

The second way in which our observation about twin desiderata may be used is to provide guidance on model building. For example, a scientist should try to select a model and a question about that model which leads to a tractable computation.

There is, of course, a long history of doing this; we often linearize nonlinear problems. What I'm looking for are novel trade-offs between models and computational difficulty.

General Remarks

R. E. Gomory

There is an old saying, "It's not what you don't know that hurts you, its what you don't know that you don't know that hurts you." If it is good to know what we don't know, it may be even better to know what things we will never know, which is what this meeting is all about.

Most people don't have much feeling for the limits of present knowledge, let alone for the permanent limits of knowledge. There is a built-in bias in education because it teaches what is known, not what isn't known. As a result we generally think more is known than actually is. This bias could be corrected if we has a more systematic characterization of the limits of present knowledge, since that characterization could be taught as known. Unfortunately we don't at present have even that,

As a result of this bias we are often unaware when we are on the edges of knowledge. Those of us who were taught the history of the Persian Wars didn't know that the events so vividly described were all from one source. There is only one for that whole period of Greek history. It may be wrong on the facts it may he biased or even made up, but that is all there is. And anything else at all that happened in Greece in that period was unknown and in all probability is unknowable. But this fragment of knowledge on the edge of the unknown, was not taught to us as that. We thought it was just a selection from a vast array of knowledge of about Greek history of that period.

There are a few commonly recognized unknowables, A person who thinks he can predict the outcome of a roulette wheel is regarded today as a crank. But persons who run Casinos work with the known averages of these unknowable individual events and so for them this is just another business. More recentlv certain aspects of weather prediction have been classed as unknowable. These are all fairly clear examples of the known and the unknowable.

Less clear are the following: People make a living picking stocks for mutual funds, Presidents run for office, on what they are going to do or have done for the economy. Which of these last two represents dealing with the unknowable, I don't know. We seem to be surrounded by actions and assertions like this dealing with the borderline unknowable. It would be real progress to know where we really stand.

APPENDIX

17–Jun–94

Einstein On the Integers

GREGORY J. CHAITIN

I've tracked down Einstein's remark that the integers are an arbitrary construct like any other, just a free invention of the human mind, justified pragmatically because they are useful. (In our discussion, I had used this remark to justify adding non-self-evident axioms to elementary number theory because of their usefulness.)

Einstein made essentially this remark in his article "Remarks on Bertrand Russell's Theory of Knowledge" in the Paul Arthur Schilpp volume *The Philosophy of Bertrand Russell,* which was Volume V in the Library of Living Philosophers, and was published by Open Court in 1944. On page 287, the English translation of Einstein's German on the facing page states:

> As a matter of fact, I am convinced that even much more is to be asserted: the concepts which arise in our thought and in our linguistic expressions are all—when viewed logically—the free creations of thought which cannot inductively be gained from sense experiences.
>
> Thus, for example, the series of integers is obviously an invention of the human mind, a self-created tool which simplifies the ordering of certain sensory experiences. But there is no way in which this concept could be made to grow, as it were, directly out of sense experiences. It is deliberately that I chose here the concept of number, because it belongs to pre-scientific thinking and because, in spite of that fact, its constructive character is still easily recognizable. The more, however, we turn to the most primitive concepts of everyday life, the more difficult it becomes amidst the mass of inveterate habits to recognize the concept as an independent creation of thinking.

Here is a related remark of Einstein's in his "Autobiographical Notes" in Paul Arthur Schilpp's, *Albert Einstein: Philosopher-Scientist,* Volume VII of the Library of Living Philosophers, Open Court, 1949, p. 13:

> All concepts, even those which are closest to experience, are from the point of view of logic freely chosen conventions . . .

Let me finally remark that it is not surprising that the integers weren't sacred to Einstein. Remember that special relativity challenged deeply held beliefs about the fundamental nature of space and time, beliefs that seemed as fundamental in physics as the integers are in mathematics.

20–Jun–94

Einstein On the Integers

PATRICK SUPPES

I want to comment on Einstein on the integers. I agree with his remark that the integers are not God given, but I take with a large grain of salt his statement that "all concepts are from the point of view of logic freely chosen conventions ... " This kind of conventionalism was also on various occasions advocated by Poincare and popularized in his books on philospy of science. But it is, in my view, hard to make stick historically or in any detailed way to be supported by experience. We can indeed hold that our theories are underdetermined by experience and even, in specific contexts, by the best experiments we can design. But there is a vast difference between being underdetermined and being arbitrary. Try publishing an arbitrary theory.

Of course, there is another reading of Einstein's remark. It is that logic determines very little. So the only criterion for logical possibility is that of freedom from contradiction. This reading I very much agree with. It is why, in the end, the program of Frege and Russell to derive mathematics from logic failed and is no longer taken seriously.

How to think about the integers is another matter. The empirical need to count is irresistible and far from being anything like a convention. The wonderful step of abstracting from counting sheep or pigs and counting jugs of wine to a general concept of natural number (but not zero) occurred in the very dim past and the details are far from clear, but it was a glorious abstraction. Already by the Hammurapi period, 1800–1600 BC, Babylonian arithmetic was quite advanced, but used for endless practical tasks. The scribes and others who learned it were almost certainly in no sense learning pure mathematics in our sense, but important routines connected with good government and estate administration.

So, looked at historically and in terms of practical use, there was not much freedom of concepts available, in relation to experience rather than logic, in the development of arithmetic over several thousand years. Against this view it might be said there was freedom in development of the particular notation and algorithms developed. There is something to this view, but each step in improvement was a small one and did not have much sense of freedom. For example, notation for zero was introduced only after about a thousand years after the Hammurabi period in Babylonian mathematics.

22–Jun–94

Non-Turing Computation

F. A. Doria

This is intended as a summary. To my knowledge the first claim of a finitary yet abstract kind of computation appears in Gödel's 1958 *Dialectica* paper, where he proves the consistency of arithmetic. (There is a 1972 English translation; both can be found in the Collected Works volumes.) Then there is Bruno Scarpellini's paper (*Zeits. Math. Logik & Grundl. der Math.*, 9, 265, 1962) where he asserts that an analog computer can (theoretically, at least) decide some undecidable predicates. Next step is Kreisel's paper ("A notion of mechanistic theory, in P. Suppes, ed., *Logic and Probability in Quantum Mechanics*, D. Reidel, 1976). Kreisel asserts that *classical* physics contains many examples of noncomputable phenomena; he bases his claims on Richardson's 1968 paper (D. Richardson, *J. Symbol. Logic*, 33, 514 (1968). Also out of Richardson's paper da Costa and myself showed in 1990 that there are infinitely many representations in several fragments of axiomatic set theory for the halting function (the function that solves the halting problem). One reference is *Found. Phys. Letters*, 4, 363 (1991). Again analog computers come to the foreground, since they can (at least in theory) decide undecidable predicates. We also showed (still unpublished) that the computation theory we obtain by adding analog stuff to ordinary Turing computation is 'equivalent' to Carnap's rule in arithmetic. Quotation marks have been added to emphasize some abuse of language here.

There are several other interesting features of that analog plus Turing computation theory that delve on complexity issues, but I would like to add a couple of lines on its physical feasibility. As Pat Suppes once remarked to me, even if its whole power is technically inaccessible, some kind of partial physical realization might be available in order to speed up in a novel way our current computers - which obviously fall short of the highly idealized Turing model.

P. S. I know that there are several other examples of beyond-Turing models, e.g. Dana Scott's. But I restricted the above summary to things with a more or less intuitive physical realization.

22–Jun–94

Einstein On the Integers

LEE A. SEGEL

Dear Greg and Patrick and colleagues, Concerning the above–I feel a bit uncomfortable with argumentae (my Latin is nonexistent, so excuse the improvised plural) ad hominem. For a solid point, we would want something that is "agreed," not just said by one person, however great .... And is "the empirical need to count irresistible." I can't trace a reference, but I seem to recall reading that certain "primitive tribes" count "one, two, many." If true this would demonstrate that indeed integers are a construct of (some) human mind(s).

25–Jun–94

Limitations and Invariance

PATRICK SUPPES

In many important cases limitations of knowledge generate principles of invariance. Here are five examples.

1. That geometric shape is not affected by orientation leads to the ancient Euclidean principle of invariance of congruence under changes of orientation.
2. Our inability to determine the absolute position in space and time of an event leads to the principle of Galilean relativity in classical physics.
3. Our inability to make an absolute separation of space and time leads to the principle of special relativity. A simple axioms that leads directly to the Lorentz transformations is the invariance of the relativistic distance between two points under a change of inertial frame of reference.
4. Our inability to distinguish continuity of physical quantities from sufficiently fine-grained discreteness leads by a modern route to Kant's Second Adntinomy and from there to a restricted isomorphism between continuous and discrete representations of many phenomena.
5. Kant's Third Antinomy concderns deterministic versus indeterministic causal theories. In the case of Sinai billiards (and others as well) there is a provable isomorphism between the classical mechanics of elastic collisions and Bernoulli flows (and for finite partitions even Bernoulli shifts). Thus there is a principle of invariance expressing our inability to distinguish by empirical observation between these two fundamental ways of thinking, ways usually thought to be in strong opposition to each other.

Finiteness and Real-World Limits

JOHN L. CASTI

A key, if not the key, element in making the leap between limits (of the Gödelian type or otherwise) in models of reality and the real system itself is to be able to spin a convincing tale as to why the model is in any meaningful way homologous, or congruent, to the system it purports to represent.

Thinking about the two sides of this picture, the real system and its model, we find an absolutely fundmental mismatch between the two. This mismatch is so fundamental, in fact, that it's usually just swept under the rug and not even acknowledged. It is simply the complete ***incongruence*** between the number systems used in the model and in the observations of the real-world process. In just about every mathematical model I've seen, it's just assumed that the variables involved take their values in some convenient number field like the real or the complex numbers. Sometimes, if the investigator is very daring, s/he may specify the integers. But in all cases, it's just assumed that there is an infinite set of numbers available. On the other hand, only a finite set of numbers has ever been observed or written down in all of human history—and that situation is not going to change anytime soon. So what we have here is as basic a distinction as one can conceive of in mathematics: the difference between the finite and the infinite.

In general, modelers simply assume that the finite set of numbers upon which their real-world system is based is big enough that, for all practical purposes, it can be considered infinite. They then merrily roll along on that assumption, using their favorite infinite number field and all that that implies and allows. Unfortunately, for studies of limits we can't proceed in such a cavalier fashion, and this finite versus infinite mismatch cannot be ignored without running the risk of throwing out the baby with the bathwater. There are at least two reasons why:

A. Finite number fields *are* the number systems of reality—at least insofar as we perceive reality by making observations and measurements.

B. There are no undecidable statements in formal systems whose statements are expressed over finite number fields, i.e., when it comes to real-world systems described by real-world number systems, there can be no Gödel's Theorem.

This second point is worth expanding upon, since it implies that when we talk about real systems and their corresponding real-world number system, every statement we can make about the system is either true or false; there can be no logically undecidable propositions about such a system. Furthermore, the statement can, in principle, be decided by carrying out a finite computation. Thus, the issue of limits to scientific knowledge about the real-world system then reduces to the question of computational tractability, e.g., the P=?NP Problem.

So my very provisional conclusion from all this is that a good place to start in a search for limits to scientific knowledge for *real,* as opposed to mathematical, systems is with setting Gödel's results to the side as being irrelevant to the matter. Instead of being distracted by this kind of mathematical red herring, we can much more productively center our attention on *finite* mathematics and the limitations that the theory of computation puts before us for this class of problems.

Finite

E. ATLEE JACKSON

I would like to add my vote of agreement for John Casti's latest message, in his 'finiteness and real-world limits.' I draw your attention to my SFI report 94–06–035, 'Evolving foundations of science,' and the distinction between the types of information contained in physical observations, math models and computer experiments—most importantly the finite duration of observations and computations done by scientists—which I also pointed out at the workshop. There are many challenging encoding/decoding questions to be unraveled between our MM and CE results and the things we observe in science - particularly in relationship to this finite duration of observations.

Also, if you are really lucky and can lay your hands on Vol. I of my books *Perspectives of Nonlinear Dynamics*, you might glance at section 4.11, and more particularly at the great piece of art in Fig. 4.76.

29–Jun–94

John Casti's page on "Finiteness and Real-World Limits"

ROLF LANDAUER

These points are an articulate statement of things I have tried to say for a long time. Of course, I make a stronger claim. I do not just question whether the real numbers, or the integers, characterize the physical world. I assert that the real numbers, or unlimited set of integers, have no available physical representation, and therefore cannot be specified, manipulated, used, etc. They do not "exist" just because they have been named and can be invoked in a closed formal axiom system.

I maintain an open mind about the existence of a largest number. I suspect that God was not that digital in orientation, and that there is no sharp cutoff, but rather a statistical sort of blur that sets in as numbers get larger.

29–Jun–94

Some Thoughts About the Continuum

Joseph Traub

This posting was stimulated by Rolf Landauer's very thoughtful message. Rolf writes "Algorithms invoked by our ordinary continuum mathematics are not executable, and therefore not suitable as a basis for the laws of physics. Algorithms for real use need to be restricted to calculations which can, at least in principle, be carried out."

I am in complete agreement with Rolf's second sentence but believe that it does not lead to the conclusion reached in his first sentence. For example, in the field of information-based complexity, see [1], we study the computational complexity of computing approximations to the solutions of continuous models. We are certainly interested in "algorithms for real use." For example, we are starting to investigate the approximate computation of path integrals.

I will consider the continuum in "three worlds":

- Phenomenon (Real World)
- Mathematical Model (Mathematical World)
- Computation (Computer World)

For more regarding these three worlds, see [2]. I'll discuss these worlds, in turn.

I'll choose physics as my example of the real world. The laws of physics must, of course, be developed by physicists. It is my understanding that only a rather small number of physicists believe that space and/or time are ultimately discrete. (This group does include some major figures such as T. D. Lee). Furthermore, if space or time are discrete it is at scales many orders of magnitude smaller than Planck's constant.

Mathematical models are conceived by physicists or applied mathematicians. There is, of course, considerable interest in discrete models such as cellular automata. However, I believe that continuous models, such as path integrals and differential equations, will continue to be important for at least the forseeable future.

I'll devote the remainder of this note to computational models. One chooses an abstraction of the computer called the "Model of Computation." I will compare and contrast two abstractions: The Turing Machine and the Real Number Model. In the Real Number Model we use the abstraction that we can store real numbers and perform exact arithmetic on them at constant cost. The Turing Machine is typically used by people working on discrete problems while the Real Number Model is used by those working on continuous problems. The latter includes people working in information-based complexity as well as those using Blum-Shub-Smale machines. See [3] for further discussion.

The choice of model in computation is important. The two models we have been discussing are not polynomially equivalent!

I will list the pros and cons of these two models of computation followed by discussion.

– Turing Machine Model –

Pro: Simple, robust

Con: Not predictive for scientific computation

– Real-Number Model –

Pro: "Natural" for continuous mathematical models; predictive for scientific computation; utilizes the power of continuous mathematics

Con: Attractive to use finite-state model for finite-state machine

I'll now discuss these pros and cons. The attraction of the Turing machine is its simplicity. Furthermore, as is often pointed out, Turing Machine results are polynomially equivalent to those for any finite model. That is, they are robust. As already observed above, results are not polynomially equivalent to those of the Real- number model.

A negative of Turing machines is that results, such as running times, are not predictive of scientific computation on real digital computers.

I'll turn to the Real-number model. I believe it is natural to match the computer model to the mathematical model. Since the reals are used in continuous mathematical models, they should also be used in the computer model.

Furthermore, the Real-number model is predictive for scientific computation. That is, modulo stability, computational complexity results in the real number model are the same as for the finite precision floating point numbers used in scientific computation.

The final pro is that one has at hand the full power of continuous mathematics. Here's one example of the significance of that. There has been considerable interest in the physics community in the result that there exist differential equations with computable initial conditions and non-computable solutions. This follows from a theorem on ill-posed problems established by Pour-El and Richards. They use computability theory to establish their result and devote a large portion of a monograph, [4], to developing the mathematical underpinnings and to proving the theorem.

An analogous result has been established using information-based complexity. The proof takes about one page; see [5]. More importantly, in information-based complexity, it is natural to consider the average case. It was recently shown that every ill-posed problem is well-posed on the average for every Gaussian measure. There is no corresponding result using computability theory.

Finally, we state a negative for the Real Number Model. Since the digital computer is a finite state machine, it is attractive to match it with a finite model of computation.

Much more can be said on this very interesting issue of the continuum but thats enough for this note.

Bibliography

[1] J. F. Traub, G. W. Wasilkowski, and H. Wozniakowski, *Information-Based Complexity,* Academic Press, 1988.

[2] J. F. Traub and H. Wozniakowski, "Theory and Applications of Information-Based Complexity," in *1990 Lectures in Complex Systems,* Santa Fe Institute. Edited by Lynn Nadel and Daniel L. Stein, Addison-Wesley, 1991.

[3] J. F. Traub and H. Wozniakowski, "Information-Based Complexity: New Questions for Mathematicians," *Mathematical Intelligencer,* 13, 34–43 (1991).

[4] M. B. Pour-El and J. I. Richards, *Computability in Analysis and Physics,* Springer Verlag, 1988.

[5] A. G. Werschulz, "What is the Complexity of Ill-Posed Problems?", *Numer. Func. Anal. Optim.,* 9, 945–967 (1987).

[6] J. F. Traub and A. G. Werschulz, "Linear Ill-Posed Problems Are Solvable on the Average for All Gaussian Measures," *Mathematical Intelligencer,* 16, 42–48 (1994).

Finiteness May Also Be a Red Herring

F. A. Doria

Let me intervene a third time in our discussion, but I think that finiteness may also be a red herring. Take for example the $P = NP$? question, which was mentioned by John Casti. The 'obvious' thing is that either we have $P = NP$ or we have $P < NP$. After all, satisfiability problems are given by nice finite Boolean expressions and so on. So, we might somehow be able to decide which alternative is correct. Yet, weird things may creep up. First is the Solovay et al. 1975 result on oracles that either given $P = NP$ or give $P < NP$. I was told that Solovay conjectures that $P = NP$ might actually be *undecidable* wrt say ZF. da Costa and myself have also speculated several funny possibilities. For instance, if one weakens the Axiom of Choice things might behave in an odd way for truly arbitrary satisfiability problems (= an infinite sequence of Boolean expressions in conjunctive normal form). Also, within a strict axiomatic framework, we may offer *two* plausible definitions for $P = NP$. Let us call one of them the 'weak' form, and the other the 'strong' form. We might end up with the following result: (Given consistency of ZF etc.), "If $P = NP$ weakly, then $P < NP$ strongly." Nice, huh? (Please, notice the following caveat: the two definitions are real stuff; the things that follow, on the effect of AC and the weird statement about $P = NP$, are reasonable hunches which we are trying to make into real, solid (finite?) stuff.)

Back to finiteness: the later writer Jorge Luis Borges once made the following comment, "the number of metaphors which we can conceive is finite." I think it easier to accept finiteness in the sense of existential philosophy than within a mathematical framework. Infinite is Cantor's paradise—or Cantor's hell. Anyway sinfully are we lured into it.

Finiteness May Also Be a Red Herring

F. A. DORIA

Let me intervene a third time in our discussion, but I think that finiteness may also be a red herring. Take for example the $P = NP$? question, which was mentioned by John Casti. The 'obvious' thing is that either we have $P = NP$ or we have $P < NP$. After all, satisfiability problems are given by nice finite Boolean expressions and so on. So, we might somehow be able to decide which alternative is correct. Yet, weird things may creep up. First is the Solovay et al. 1975 result on oracles that either given $P = NP$ or give $P < NP$. I was told that Solovay conjectures that $P = NP$ might actually be *undecidable* wrt say ZF. da Costa and myself have also speculated several funny possibilities. For instance, if one weakens the Axiom of Choice things might behave in an odd way for truly arbitrary satisfiability problems (= an infinite sequence of Boolean expressions in conjunctive normal form). Also, within a strict axiomatic framework, we may offer *two* plausible definitions for $P = NP$. Let us call one of them the 'weak' form, and the other the 'strong' form. We might end up with the following result: (Given consistency of ZF etc.), "If $P = NP$ weakly, then $P < NP$ strongly." Nice, huh? (Please, notice the following caveat: the two definitions are real stuff; the things that follow, on the effect of AC and the weird statement about $P = NP$, are reasonable hunches which we are trying to make into real, solid (finite?) stuff.)

Back to finiteness: the later writer Jorge Luis Borges once made the following comment, "the number of metaphors which we can conceive is finite." I think it easier to accept finiteness in the sense of existential philosophy than within a mathematical framework. Infinite is Cantor's paradise—or Cantor's hell. Anyway sinfully are we lured into it.

01–Jul–94

Advertisement For a Paper I Like

ROLF LANDAUER

All of us have papers which reached us when we were very young and influencable, and which had a lasting effect. And when we later on tell others about one of those, we often find they do not share our enthusiasm. Let me try anyway, quoting from one of my semi-published papers.

My personal story starts with P. W. Bridgman, Nobel laureate, high-pressure physicist, and author of a number of discussions of "operationalism." The latter can be considered to be a Yankee variant of logical positivism. Bridgman, in 1934, published a paper "A Physicist's Second Reaction to Mengenlehre," in Vol. II of *Scripta Mathematica* (p.3). This journal was published at Yeshiva University and oriented to fundamental questions in mathematics. The paper starts with a remarkable qualifying note by the editor: "The Scripta Mathematica' takes pleasure in publishing this interesting article by Professor P. W. Bridgman. As in the case of all articles, however, it is understood that this does not necessarily represent the views of the editors." Bridgman was, by that time, already well known, and science is normally full of divergent viewpoints. Fifty years later, we can only guess that the unusual editorial qualification represented the editor's desperation. He may have been hesitant to turn down a famous scientist and, at the same time, was afraid of the wrath of his readers. Bridgman's paper presents his attempt to wrestle with the paradoxes of set theory. He objects to self-referential definitions and asks for mathematical procedures in which each step can be carried out directly and unambiguously. In a remarkable coincidence(?) he uses the word "program" for such a succession of instructions. Bridgman would have been entirely satisfied by a restriction of mathematics to that which can be done on a Turing machine, which came along a few years later. Bridgman did not go on to ask about the "physical realizability" of these operations.

P. S. Papers by famous people often remain obscure. Only new Ph.Ds believe that they are the only people whose work is neglected. A 1917 paper by Einstein took many decades to become widely known. Despite the fact that it was delivered to the German Physical Society the year he retired as its President.

04–Jul–94

More About Finiteness

F. A. DORIA

I pondered it throughout the evening. After reading Joe Traub's one-pager today I think I can summarize my own discomfort in two questions:

1. Mathematics is a coarse, unsubtle language. Nevertheless, what is "finite" in mathematics?
2. What can we do in an ordinary language?

1. Mathematics is coarse. Think about the following assertions: "All men are mortal." "Every man is mortal." "Each man is mortal." There are subtle differences among them, yet we (*usually*) formalize all three in, say, $\forall x(x \in \text{Men} \rightarrow x \in \text{Mortal}$.

Finiteness. When is a set "finite"? There are at least two reasonable definitions for it, and they are not equivalent:

Standard finiteness: A set is finite when there is a natural number n which is the cardinal number of that set.

Dedekind finiteness: A set is finite when none of its proper subsets can be placed in 1-1 correspondence with the whole set. (Actually the definition goes back to Galileo.)

These are not equivalent: if the axiom of choice fails, we can find a counterexample. Try, say, talking about "smallness" instead of "finiteness." We may sketch the following: $\emptyset$ is a small set. If $A \subset B$, and if B is small, so is A. If A is small, then $A \cap B$ is small. We have an ideal of sets. "Large" sets are reasonably described by filters (the dual ones). Now, recall that the large cardinal stuff which is in most cases undecidable within ZFC is related to certain filters over segments of real numbers. Our intuitions are rapidly lost here.

2. Can we talk about finiteness, decidability, intractability in an ordinary language? To be frank, I don't know. Perhaps one might somehow discover a sort of middle road, a much finer kind of mathematics where our intuition follows us deeper into the workings of the formalism, and where Joe's questions wouldn't sound as reasonable as they sound today.

So far, I stick to formalism, where I can stand on firm ground. A firm ground that anyway might dissolve into a boundless *sunya*. Shanti shanti shanti.

07–Jul–94

Alternative Realities

JOHN L. CASTI

In our deliberations on Limits, there are three different worlds to consider. These are the

- **Real World,** consisting of actual, i.e., *real,* observations and measurements;

- **Mathematical World,** in which live abstractions like real numbers, continuity, randomness and all the other taken-for-granted fictions that we regularly employ in painting pictures of reality;

- **Computational World,** where algorithms are executed and computations take place—on either real or abstract computing devices.

On the basis of what's been sent to the Limits BBS over the past month or so, I think it's worth examining each of these worlds in a lot more detail than has heretofore been the case, in an attempt to sort out and pin down more precisely some of the questions that people are asking in our Limits discussions. In particular, my feeling is that we have to be much more aware of the number systems available for use in these three worlds. This is primarily because for whatever psychological reasons, humans seem to be predisposed to use numbers to encode our observations and descriptions of each of the foregoing worlds. Thus, the character of what we can and cannot know—or at least *say*—about questions posed in these worlds generally rest on the assumptions we make about the kinds of number systems at our disposal to describe phenomena of interest.

– Real Worlds –

The key word describing these worlds of nature and human affairs is *finiteness;* in any real world, everything is finite—measurements, lengths and durations. As a result, the only kinds of number systems we need to describe what's going on in any real world are finite systems, e.g., finite fields like Z mod p (p a prime) or even more primitive finite sets like 0, 1, 2, . . . , 9.

I have heard arguments to the effect that while it is certainly true that in practice we can only measure things in the real world to finite precision, any particular measurement could be *in principle* any real number. Let me offer several arguments against this line of reasoning:

1. This real-world-measurement context is just one of many in which the phrase "in principle" is completely meaningless. I would defy anyone to give a *physically realizable* procedure for measuring any real-world quantity that results in a real number. Furthermore, it is just not the case that the measurement of a physical quantity such as the length of the side of a room, or for that matter, the size of the universe can "in principle" be any real number. Since every physical quantity is finite, it necessarily follows that every such quantity is bounded; ergo, a measurement of that quantity is also bounded, which means that its value *cannot* be any real number—even "in principle."

2. It seems to be part of the folk wisdom of the scientist's, and especially the physicist's, world view that assuming real numbers makes life easier somehow. It's hard to quarrel with three centuries of successes in physics that is to a great extent based on this assumption, although I might want to qualify the word "success." But that's a topic for a whole new paper. So let me content myself here with simply pointing out that this assumption is equivalent to assuming an axiomatic framework for the measurement process that includes the Axiom of Choice, one of the most notorious logical pitfalls in all of mathematics. One can only wonder whether the benefits from this Axiom are worth the costs in potential paradoxes stemming from this Axiom (e.g., Banach-Tarski) that fly flat in the face of physical common sense.

   In this regard, it's worth noting that assuming the reals, or equivalently the Axiom of Choice, is certainly an act that would raise the eyebrows of William of Ockham, since it's a completely unnecessary assumption. As per point #1, you don't need an infinite set of numbers to describe the physical world, however convenient it might be for some mathematical purposes. So why trot out the blunderbuss of the reals when all you need to bag your quarry is the BB gun of a finite set?

   Note that I'm **not** claiming that something like the Axiom of Choice necessarily leads to paradoxes in the physical world. Who knows if use of the Axiom of Choice will lead to reality paradoxes or not. All we can say for sure is that it leads to things like the Banach-Tarski paradox, which certainly does run counter to our physical intuition about the way of the world, as it seems to violate conservation of mass, among other things. (Just for the record, the B-T Paradox shows how you can cut up a sphere, say, into a finite number of pieces, and then re-assemble these pieces to form a new sphere whose radius is double that of the original.) So the Axiom of Choice *might* lead to a paradox, and I wonder why we should run such a completely unnecessary
   risk?

3. For a physicist, assuming an infinite set of numbers has its mathematical attractions. But we're not talking about mathematical abstractions when we're situated in the world of real phenomena; we're talking about actual observations and measurements. And when it comes to asking if there are unanswerable questions relating to real-life, get-your-hands-dirty observations and measurements, it's a positive impediment to be saddled with something like the real numbers. There are many reasons for this, but the one that I feel is most germane to our Limits deliberations is simply that all statements made in a framework involving finite number fields are decidable; there cannot exist a Gödel-like theorem in this setting. Therefore, by restricting ourselves to finite sets of numbers for our measurements, we can safely ignore Goedelian notions of undecidability in our study of questions of unanswerability, which all then reduce to the problem of computational intractability.

– Mathematical Worlds –

These seem to be the only worlds in which just about all of us in the Limits teleconference can seem to agree, probably because in these worlds anything that's not logically forbidden eventually becomes necessary. So if you want continuity, infinite sets, random numbers, inherent stochasticity, times going to infinity or anything else, just snap your fingers, and presto! you can have it, as if by magic. Perhaps this is why so many of the Limits postings have been about one or another of these worlds: they're easy, since you can make them up as you go along, dispensing whatever properties you like as conditions, needs and wishes dictate. Anyway, they are not the focus of this missive, so I'll pass over them in silence.

– Computational Worlds –

The worlds of computation occupy the curious position of having a foot in both the real and the mathematical worlds. When we talk about carrying out an actual calculation on a physical device (e.g., fingers, abacus, slide rule or computing machine (digital or analog)), we're talking about the real world. And in that case, all the restrictions noted above apply to any such computations. On the other hand, when we talk about a *model* of computation (e.g., a Turing machine or a BSS machine), then both our feet are planted firmly in the world of mathematics. Such models are mathematical theories of computation, and as such they may or may not bear any resemblance to *real-world* computational processes.

What this means, of course, is that the kind of restrictions on what can be done with, say, a Turing machine, may or may not be relevant to what can be done with a Connection Machine 5 or a SPARC workstation. And, as noted above, it is the limitations on *real* computation that ultimately dictate the questions that we can or cannot answer about the *real* world. Passing through any of the mathematical worlds—including those consisting of models of computation—can be at best only a waystation on the journey to understanding what we can possibly know about a given real-world process.

– In Conclusion –

I strongly suspect that there's very little in this didactic *cri de coeur* that we're not all intimately familiar with, at least implicitly. Nevertheless, I hope that the points raised above help focus our attention more *explicitly* on some of the issues that we'll have to grapple with in coming to terms with Limits. Each of the three classes of worlds—real, mathematical, computational—has its own characteristic set of problems, and it's certainly not my intention to suggest here that any particular set is unworthy of serious investigation. Quite the contrary, in fact. But it IS my intention to suggest that: (1) we keep uppermost in mind the particular world we're talking about when we make pronouncements about what can and cannot be known in a given area, and (2) we focus a disproportionate amount of our energies on the *interfaces* between these very different worlds. To my eye, there has been way too much implicit wishful thinking in many of the Limits postings, to the effect that work in the worlds of mathematics and/or theoretical computation has something relevant to say about what's happening in what I called above the Real World. Maybe (even probably) it does. But if so, I'd like to see an explicit chain of reasoning establishing that fit. I'm sorry to say that it's just not good enough to *believe* or *feel* that there's a connection; you have to actually *prove* it!

So, in my view the tone of our deliberations will be greatly elevated by making a more careful distinction between the map and the territory, which, in turn, will lead to an on-going interchange generating at least a bit more light than heat.