

Church's thesis: computability, proof, and open-texture

Stewart Shapiro

Church [1936]

We now define the notion . . . of an *effectively calculable* function of positive integers by identifying it with the notion of a recursive function of positive integers . . . This definition is thought to be justified by the considerations which follow, so far as positive justification can ever be obtained for the selection of a formal definition to correspond to an intuitive one. (Church [1936, §7])

Shapiro [1981, 353-354]

Computability is a property related to either human abilities or mechanical devices, both of which are at least prima facie non-mathematical. It is therefore widely agreed that the question of Church's thesis is not a mathematical question, such as the Goldbach conjecture . . . That is to say, mathematicians do not seek to show either that CT follows from accepted laws of number theory or that it contradicts such laws. Nevertheless, both mathematicians and philosophers have offered various non-mathematical arguments either for or against the thesis. Goldbach's conjecture can be settled, if at all, only by mathematical argument, but CT can be settled, if at all, only by arguments that are, at least in part, philosophical. (Shapiro [1981, 353-354])

Kleene [1971, 317, 318-319]:

Since our original notion of effective calculability of a function (or of effective decidability of a predicate) is a somewhat vague intuitive one, [CT] cannot be proved . . .

While we cannot prove Church's thesis, since its role is to delimit precisely an hitherto vaguely conceived totality, we require evidence that it cannot conflict with the intuitive notion which it is supposed to complete; i.e., we require evidence that every particular function which our intuitive notion would authenticate as effectively calculable is . . . recursive. The thesis may be considered a hypothesis about the intuitive notion of effective calculability, or a mathematical definition of effective calculability; in the latter case, the evidence is required to give the theory based on the definition the intended significance.

The following appears in a letter that Church wrote to Kleene:

In regard to Gödel and the notions of recursiveness and effective calculability, the history is the following. In discussion with him . . . it developed that there is no good definition of effective calculability. My proposal that lambda-definability be taken as a definition of it he regarded as thoroughly unsatisfactory. I replied that if he would propose any definition of effective calculability which seemed even partially satisfactory I would undertake to prove that it was included in lambda-definability.

Church's letter continues:

His [Gödel's] only idea at the time was that it might be possible, in terms of effective calculability as an undefined term, to state a set of axioms which would embody the generally accepted properties of this notion, and to do something on that basis.

Gödel [1946]

Tarski has stressed in his lecture (and I think justly) the great importance of the concept of general recursiveness (or Turing's computability). It seems to me that this importance is largely due to the fact that

with this concept one has for the first time succeeded in giving an absolute definition of an interesting epistemological notion . . .

Sieg [2002]:

My strategy . . . is to bypass theses altogether and avoid the fruitless discussion of their (un-)provability. This can be achieved by *conceptual analysis*, i.e., by sharpening the informal notion, formulating its general features axiomatically, and investigating the axiomatic framework . . .

The detailed conceptual analysis of effective calculability yields rigorous characterizations that dispense with theses, reveal human and machine calculability as axiomatically given mathematical concepts, and allow their systematic reduction to Turing computability. (Sieg [2002, 390, 391])

One formulation of proofhood has it that a proof is a derivation in Zermelo-Fraenkel set theory (ZF), or a sequence of statements that can be "translated" into a derivation in ZF. Call this a *ZF-proof*.

In a collection of notes entitled "What does a mathematical proof prove?" (published posthumously in [1978, 61-69]), Imre Lakatos makes a distinction between the pre-formal development, the formal development, and the post-formal development of a branch of mathematics. Lakatos observes that even after a branch of mathematics has been successfully formalized, there are residual questions concerning the relationship between the formal deductive system (or the definitions in ZF) and the original, pre-formal mathematical ideas. How can we be sure that the formal system accurately reflects the original mathematical structures? These questions cannot be settled with a derivation in a further formal deductive system, not without begging the question or starting a regress—there would be new, but similar, questions concerning the new deductive system.

My two earlier papers (Shapiro [1981, [1993]) and Mendelson [1990] present a few historical situations which are like CT in the relevant respects, in that they identify an intuitive, pre-theoretic notion with a more precisely defined one. These other "theses" are not subject to doubt anymore, nor is their status as mathematics in question. One especially compelling example, which we may call "Weierstrass's thesis" identifies the pre-theoretic or intuitive notion of a continuous function with the one yielded by the now standard ϵ - δ definition. It is clear that there is an intuitive such notion, and that mathematicians worked with it well before the rigorous definition was proposed and accepted. Moreover, the definition is accepted, without opposition, despite that fact that it has some consequences that conflict with intuition. Prominent among those is the existence of a continuous curve that is nowhere differentiable.

Ancient Greek mathematicians did not identify magnitudes like lengths, areas, and volumes with real numbers (as we do today, thanks to analytic geometry). The following might be called "Eudoxus's thesis":

Let x, X be two like magnitudes (i.e., two line segments, two areas, or two volumes), and let y, Y be two like magnitudes. Then $x:X = y:Y$ if and only if for any pair of natural numbers m, n , mx is longer than (resp., shorter than, resp. identical to) nX if and only if my is longer than (resp., shorter than, resp. identical to) nY .

Waismann introduces the notion of open-texture in an attack on crude phenomenalism, the view that one can understand any cognitively significant statement in terms of sense-data. The failure of this program:

is not, as has been suggested, due to the poverty of our language which lacks the vocabulary for describing all the minute details of sense experience, nor is it due to the difficulties inherent in producing an *infinite* combination of sense-datum statements, though all these things may contribute to it. In the main it is due to a factor which, though it is very important and really quite obvious, has to my knowledge never been noticed—to the "open texture" of most of our empirical concepts. (Waismann [1968, 118-119])

Here is one of the thought experiments that Waismann uses to characterize the notion of open-texture:

Suppose I have to verify a statement such as ‘There is a cat next door’; suppose I go over to the next room, open the door, look into it and actually see a cat. Is this enough to prove my statement? . . . What . . . should I say when that creature later on grew to a gigantic size? Or if it showed some queer behavior usually not to be found with cats, say, if, under certain conditions it could be revived from death whereas normal cats could not? Shall I, in such a case, say that a new species has come into being? Or that it was a cat with extraordinary properties? . . . The fact that in many cases there is no such thing as a conclusive verification is connected to the fact that most of our empirical concepts are not delimited in all possible directions.

The notion of gold seems to be defined with absolute precision, say by the spectrum of gold with its characteristic lines. Now what would you say if a substance was discovered that looked like gold, satisfied all the chemical tests for gold, whilst it emitted a new sort of radiation? ‘But such things do not happen.’ Quite so; but they *might* happen, and that is enough to show that we can never exclude altogether the possibility of some unforeseen situation arising in which we shall have to modify our definition. Try as we may, no concept is limited in such a way that there is no room for any doubt. We introduce a concept and limit it in *some* directions; for instance we define gold in contrast to some other metals such as alloys. This suffices for our present needs, and we do not probe any farther. We tend to *overlook* the fact that there are always other directions in which the concept has not been defined. . . . we could easily imagine conditions which would necessitate new limitations. In short, it is not possible to define a concept like gold with absolute precision; i.e., in such a way that every nook and cranny is blocked against entry of doubt. That is what is meant by the open texture of a concept. (Waismann [1968, 120])

The phrase “open texture” does not appear in Waismann’s treatment of the analytic-synthetic distinction in a lengthy article published serially in *Analysis* ([1949], [1950], [1951], [1951a], [1952], [1953]), but the notion clearly plays a central role there. He observes that language is an evolving phenomenon. As new situations are encountered, and as new scientific theories develop, the extensions of various predicates change. Sometimes the predicates become sharper and, importantly, sometimes the boundaries move. When this happens, there is often no need to decide, on hard metaphysical or semantic grounds, whether the application of a given predicate to a new case represents a change in its meaning or a discovery concerning the term’s old meaning. And even if we focus on a given period of time, language use is not univocal:

Simply . . . to refer to “the” ordinary use [of a term] is naive. There are *uses*, differing from one another in many ways, e.g. according to geography, taste, social standing, special purpose to be served and so forth. This has long been recognized by linguists . . . [These] particular ways of using language loosely [revolve] around a—not too clearly defined—central body, the standard speech . . . [O]ne may . . . speak of a *prevailing* use of language, a use, however, which by degrees shades into less established ones. And what is right, appropriate, in the one may be slightly wrong, wrong, or out of place in others. And this whole picture is in a state of flux. One must indeed be blind not to see that there is something unsettled about language; that it is a living and growing thing, adapting itself to new sorts of situations, groping for new sorts of expression, forever changing. (Waismann [1951a, 122-123])

In the final installment of the series, he writes: “What lies at the root of this is something of great significance, the fact, namely, that language is never complete for the expression of all ideas, on the contrary, that it has an essential *openness*” (Waismann [1953, 81-82]).

Waismann points out that major advances in science sometimes—indeed usually—demand revisions in the use of common terms: “breaking away from the norm is sometimes the *only* way of making oneself understood” ([1953, 84]). If I may interject my own example, we can wonder if scientists contradicted themselves then they said that atoms have parts. Well, what is it to be an atom? Surely at some point in history, “atom” could have been defined to be a particle of matter that has no proper parts. Metaphysicians still use the term that way.

One might think that in cases like these, a new word, with a new meaning, is coined with the same spelling as an old word, or one might think that an old notion has found new applications. Did Einstein discover a hidden and previously unnoticed relativity in the established meaning of the word “simultaneous”? Or did he coin a new term, to replace the old, scientifically misleading one? According to Waismann, there is often no need, and no reason, to decide what counts as a change in meaning and what counts as the extension of an old meaning to new cases—going on as before as Wittgenstein might put it. As Waismann said, in an earlier installment: “there are no *precise rules* governing the use of words like ‘time’, ‘pain’, etc., and that consequently to speak of the ‘meaning’ of a word, and to ask whether it has, or has not changed in meaning, is to operate with too blurred an expression” (Waismann [1951, 53]).

Before returning to computability, let us briefly re-examine Lakatos’s favorite example, developed in delightful detail in *Proofs and refutations* [1976]. He presents a lively dialogue involving a class of rather exceptional mathematics students. The dialogue is a rational reconstruction of the history of what may be called Euler’s theorem:

Consider any polyhedron. Let V be the number of vertices, E the number of edges, and F the number of faces. Then $V-E+F = 2$.

The dialogue opens with the teacher presenting a proof of Euler’s theorem. We are told to think of a given polyhedron as hollow, with its surface made of thin rubber. Remove one face and stretch the remaining figure onto a flat wall. Then add lines to triangulate all of the polygonal faces, noting that in doing so we do not change $V-E+F$. For example, drawing a line between two vertices of the same polygon adds 1 face and 1 edge, no vertices. When the figure is fully triangulated, start removing the lines one or two at a time, doing so in a way that does not alter $V-E+F$. For example, if we remove two lines and the included vertex from the boundary, we decrease V by 1, E by 2, and F by 1. At the end, we are left with a single triangle, which, of course, has 3 vertices, 3 edges, and 1 face. So for that figure, $V-E+F = 1$. If we add back the face we removed at the start, we find that $V-E+F = 2$ for the original polyhedron. QED.

The class then considers a barrage of counterexamples to Euler’s conjecture. Examples include a picture frame, a cube with a cube-sized hole in one of its faces, a cube with cube-sized hole in its interior, and a “star polyhedron” whose faces protrude from each other in space. One of the students even proposed that a sphere and a torus each qualifies as a polyhedron, and thus counterexamples to Euler’s theorem. Since a sphere and a torus have no vertices or edges, and a single face, $V-E+F = 1$.

A careful examination shows that each counterexample violates (or falsifies) one of the three main “lemmas” of the teacher’s proof. In some cases, the figure in question cannot be stretched flat onto a surface after the removal of a face. In other cases, the stretched figure cannot be triangulated either at all or without changing the value of $V-E+F$, and in still other cases, the triangulated figure cannot be decomposed without altering the value of $V-E+F$.

Some students declare that the counterexamples are “monsters” and do not refute Euler’s theorem. One route is to insist that the figures in question are not really polyhedra. A philosopher in the crowd might argue that a meaning-analysis of the word “polyhedron” would reveal this, and then the class could get entangled in a debate over the meaning of this word in ordinary language (be it English, Greek, Latin, etc.). The class briefly considers—and dismisses—a desperate attempt along those lines: one *defines* a polyhedron to be a figure that can be stretched onto a surface once a face is removed, and then triangulated and decomposed in a certain way. That would make the teacher’s “proof” into a stipulative definition. A second maneuver is to overly restrict the theorem so that the proof holds. One group declares that the proper theorem is that for any convex, “simple” polyhedron, $V-E+F = 2$. This group is content to ignore the interesting fact that $V-E+F = 2$ does hold for some concave, and some non-simple polyhedra. A third line is to take the counterexamples to refute Euler’s theorem, and to declare that the notion of “polyhedron” is too complex and unordered for decent mathematical treatment. They just lose interest in the notion. A fourth line accepts the counterexamples as refuting Euler’s theorem, and looks for a generalization that covers the Eulerian and non-Eulerian polyhedra in a single theorem.

At the end, a most advanced student proposes a purely set-theoretic definition of “polyhedron”. Accordingly, a polyhedron just is a set of things called “vertices”, “edges”, and “faces” that satisfy some given formal conditions. It really does not matter what the “vertices”, “edges”, and “faces” are, so long as the stated conditions are satisfied. That is, the theorem has been removed from the topic of space altogether. The student in question gives a fully formal (or at least easily formalizable) proof of a generalization of Euler’s theorem from these definitions. The only residual question left, it seems to me, is the extent to which the set-theoretic definition captures the essence of the original, pre-theoretic (or at least pre-formal) concept of polyhedron.

It is straightforward to interpret the situation in Lakatos’s dialogue—or, better, the history it reconstructs—in terms of Waismann’s account of language. The start of the dialogue refers to a period in which the notion of polyhedron had an established use in the mathematical community (or communities). Theorems about polyhedra go back to ancient Greece. The mathematicians of the time, and previous generations of mathematicians, were working with a notion governed more by a Wittgensteinian family resemblance than by a rigorous definition that determines every case one way or the other. In other words, the notion of polyhedron exhibited what Waismann calls open-texture. It was not determinate whether a picture frame counts as a polyhedron. Ditto for a cube with a cube-shaped hole in one of the faces, etc. When the case did come up, and threatened to undermine a lovely generalization discovered by the great Euler, a decision had to be made. As Lakatos shows, different decisions were made, or at least proposed.

Getting back to the matter at hand, it is part of the standard argument for the received view that Church’s thesis cannot be determinately true, much less subject to proof, since computability is a vague notion, while recursiveness is sharp. One can perhaps retort that computability is itself sharp. Mendelson [1990, 232] takes the opposite retort:

The concepts and assumptions that support the notion of partial-recursive function are, in an essential way, no less vague and imprecise than the notion of effectively computable function; the former are just more familiar and are part of a respectable theory with connections to other parts of logic and mathematics . . . Functions are defined in terms of sets, but the concept of set is no clearer than that of function . . . Tarski’s definition of truth is formulated in set-theoretic terms, but the notion of set is no clearer than the that of truth.

Let us now examine Church’s thesis, and the impact of Turing’s [1936] analysis (as well as those of Gandy [1988] and Sieg [2002], [2002a]), in terms of the open-texture of mathematical concepts. Our main question, of course, concerns the so-called “intuitive”, or “pre-theoretic” notion of computability.

Typically, the extension of the modal construction underlying the suffix is sensitive to the interests of those speaking or writing, and to their background assumptions. Say that a distance and time is “runable” by me if I can cover the distance in the given time on relatively flat ground. Is an eight-minute mile (or five minute kilometer) runable by me? It depends on what, exactly, is being asked. If I warm up right now and go out and run a mile as fast as I can, it will take me much longer than eight minutes, probably ten or eleven (if I don’t get injured for overdoing it). So, in that sense, an eight minute mile is not runable by me. On the other hand, if I were to spend six months on a training regimen, which involves working out diligently with a trainer four or five days each week and losing about 25 pounds (and avoiding injury), then I probably could manage an eight minute mile again. And I am probably capable of executing this regimen. So in a sense, an eight minute mile is runable after all. If I underwent surgery and spent several years on training, and perhaps replaced some of my body parts, I might get it down to seven minutes. So, in some sense, a seven minute mile is runable. In short, it all depends on what one means by “able”.

So what are the corresponding parameters of the pre-theoretic computability? What tools and limitations are involved? I would suggest that in the thirties, and probably for some time afterward, this notion was subject to open-texture. The concept was not delineated with enough precision to decide every possible consideration concerning tools and limitations. And just as with Lakatos’s example of a polyhedron, the mathematical work, notably Turing’s [1936] argument and the efforts of the other founders—people like Church, Turing, Emil Post, Kleene, and Péter—*sharpened* the notion to what we have today. In other words, the mathematical work served to set the parameters of the original pre-theoretic notion.

The philosophical issues are illustrated in the so-called “easy” half of Church’s thesis: every recursive function is computable. This raises the more or less standard matter of idealization. The following double recursion defines an Ackermann function f (in terms of the successor function s):

$$\begin{aligned} \forall y(fy0=s0) \\ \forall x(f0sx=syf0x) \\ \forall y\forall x(fsx=fxfyx) \end{aligned}$$

The defined function is recursive, but not primitive recursive. Boolos [1987] points out that the value of $f4,4$ is larger than the number of particles in the known universe. I dare anyone to compute it. In a real sense, the Ackermann function *cannot* be computed, in much the same sense that a two minute mile is not runable (by me or anyone else).

Of course, this instance is standard, and so is the reply. Mendelson writes:

Human computability is not the same as effective computability. A function is considered effectively computable if its value can be computed in an effective way in a finite number of steps, but there is no bound on the number of steps required for any given computation. Thus, the fact that there are effectively computable functions which may not be humanly computable has nothing to do with Church’s thesis.

Note that even though we idealize, we do insist that the instructions given to the computist be finite. Otherwise, every number-theoretic function would be computable. Similarly, every function can be “computed” with a Turing machine with infinitely many states. So there are limits to the idealization.

The fact that we are idealizing on time, attention span, and materials does not, by itself, sharpen the notion all the way to the hard rails of recursiveness. Some open-texture remains, depending on how far the idealization is to go. It seems reasonable, for example, to set *some fixed* bound on how fast a function can grow before we will call it “computable”. We might require a function to be computable in polynomial time, or exponential time, or hyper-exponential time, or in polynomial space, or whatever. This would disqualify the Ackermann functions, and perhaps with good reason.

We noted that we are interested in (idealized) human computability. But what do we count as part of the computist? Even an idealized person, considered in isolation, is more like a finite state machine than a Turing machine.

In his later article Mendelson [1990, 233] observes that there is little doubt that the so-called “east half” of CT is established beyond all doubt:

The so-called initial functions are clearly . . . computable; we can describe simple procedures to compute them. Moreover, the operations of substitution and recursion and the least-number operator lead from . . . computable functions to . . . computable functions. In each case, we can describe procedures that will compute the new functions.

Mendelson concludes that this “simple argument is as clear a proof as I have seen in mathematics, and it is a proof in spite of the fact that it involves the intuitive notion of . . . computability”. I quite agree.

I submit that both conceptually and in the historical context, the proof in question serves to *sharpen* the intuitive notion. The reasoning allows us to see where the idealization from actual human or machine abilities comes in. By examining the argument, we see that we are to ignore, or reject, the possibility of a computation failing because the computist runs out of memory or materials. We thus declare the Ackermann functions to be computable, for example. The idealization away from considerations of feasibility, bounded states, and the like, is what Lakatos might call a “hidden lemma” of the proof. There is no worry about the proof allowing a sorites series to a function not computable due to feasibility. Logically, however, it would not be amiss for someone to invoke the Ackermann

function as a Lakatos-style refutation, calling for a restricted theorem, or monster-barring, or monster-adjustment, or the like. And someone else could retort to this that issues of feasibility and limits on space and memory were never part of the pre-theoretic notion. A third person could take the situation as impetus to generalize the notion of computability, yielding notions like finite state computability, push-down computability, polynomial space computability, etc. With Waismann, I do not see a strong need to adjudicate disputes concerning which of these gives the true essence of the pre-theoretic notion.

Let us now turn to the so-called harder and, at one time, controversial direction of Church's thesis, the statement that every computable function is recursive. Here, too, we ponder exactly what tools and abilities we allot to our idealized human computist. We are asking about *deterministic* computability. The computist should not act randomly, willy nilly, at any point during the computation. She is supposed to execute a fixed algorithm. OK, then what is an algorithm? Presumably, an algorithm is a set of instructions that tells the computist what to do at every stage. Is that notion sufficiently sharp, pre-theoretically?

Suppose that some idealized counterpart to a person has an intuitive ability to detect whether a given sentence in the language of arithmetic is true. Suppose that we give this person the following instruction: if the string before you is a truth of arithmetic, then output 1; otherwise output 0. Does this count as an algorithm, for our super idiot savant? After all, the instructions tell him what to do at every stage, and by hypothesis, he is capable of executing this instruction. If this does count as an algorithm, then Church's thesis is false, since the computed function is not recursive.

Clearly, this does not count as an algorithm. But what in the *pre-theoretic* notion precludes it? Presumably, the computist is not supposed to use any intuition—other than what is needed to recognize symbols. Do we have a clear, unambiguous concept of what counts as (allowable) intuition? To be sure, actual humans do not have the ability to recognize arithmetic truth. But what of our idealized computists?

Kalmár [1959, 72] gives a “plausibility argument” *against* CT. Let f be a two-place function from natural numbers to natural numbers. Define the *improper minimalization* of f to be the function f' such that:

$$f'(x) = \begin{cases} \text{the least natural number } y \text{ such that } f(x,y)=0, & \text{if there is such a } y \\ 0, & \text{if there is no natural number } y \text{ such that } f(x,y)=0. \end{cases}$$

There are recursive functions φ whose improper minimalization ψ is not recursive. Indeed, the halting problem is a improper minimalization of the function whose value is 0 if y is the code of a complete computation of the Turing machine whose code is x , started with x as input; and whose value is 1 otherwise.

Kalmár proposes the following “method” to “calculate the value $\psi(p)$ in a finite number of steps”:

Calculate in succession the values $\varphi(p,0)$, $\varphi(p,1)$, $\varphi(p,2)$, . . . and simultaneously try to prove by all correct means that none of them equals 0, until we find either a (least) natural number q for which $\varphi(p,q)=0$ or a proof of the proposition stating that no natural number y with $\varphi(p,y)=0$ exists; and consider in the first case this q , in the second case 0 as result of the calculation. (Kalmár [1959, 76-77])

I am not suggesting that there is a legitimate sense in which Kalmár's method counts as an algorithm. Even at the time, Mendelson [1963, §3] had no trouble dismissing the example. But is there something unambiguous in the *pre-theoretic* notion, or notions, of computability, in use in the 30's and a few decades after, that rules it out, definitively? The question here is why Kalmár thought that this “method” constitutes an algorithm that is relevant to Church's thesis. To say the least, he was an intelligent mathematician, and was not prone to deny what is obvious, a mere matter of understanding the meaning of a word in use.

Notice that Kalmár's instructions *do* tell the computist what to do at each stage, provided that she is a competent mathematician. She is told to try to prove a certain theorem. Mathematicians know how to do that. Entrance and

qualifying examinations test prospective mathematicians for their ability to prove theorems. Those that display this ability are admitted to the profession. If Kalmár's “plausible” assumption is correct, then his “method” will indeed terminate with the correct answer, if the computist is diligent enough and has unlimited time and materials at her disposal.

In Chapter 19 of her landmark textbook, Péter [1957, §19.2] at least tentatively endorses Kalmár's “plausibility” argument against Church's thesis. The next chapter of the book turns to Church's thesis. It opens:

Now I should like to quote some of the arguments used in attempts to make plausible the identification of the “calculable functions” with the general recursive functions.

The assertion that the values of a function are everywhere calculable in a finite number of steps has meaning only under the condition that this calculation does not depend on some individual arbitrariness but constitutes a procedure capable of being repeated and communicated to other people at any time. Hence it must be a mechanical procedure, and thus one can imagine, in principle, a machine able to carry through the single steps of the calculation.

One would think that Kalmár's “method” is thereby disqualified. Different mathematicians will proceed differently to the instruction: “try to prove such and such a theorem”, and no one has argued that this “method” can be mechanized. Indeed, it can't be. Péter goes on to give a detailed, painstaking analysis of computation, relating the notion to Turing machines. It is clearly in the spirit of Turing [1936]. The chapter closes thus:

If we assume that in the concept of a function calculable by the aid of a Turing machine we have succeeded in capturing the concept of the most general number-theoretic function whose values are everywhere calculable, then the results obtained above in fact characterize the general recursive functions as the functions calculable in the most general sense; and under this interpretation, the function [defined by Kalmár] . . . is an example of a function not calculable in the most general sense. (Péter [1957, §20.13])

So far, so good. But Péter still hedges:

But, however plausible it may seem that this interpretation correctly reflects real mathematical activity, we are nevertheless dealing here with a certain demarcation of the concept of calculability, and the future evolution of mathematics may bring about methods of calculation completely unexpected nowadays.

That was then, this is now. It seems to me that in the ensuing decades, the community of logicians has come to see the notion of computability as sufficiently sharpened. It is now reasonable to hold that Church's thesis is established with as much rigor as anything in (informal) mathematics. Gandy and Sieg have fulfilled Gödel's suggestion above, “to state a set of axioms which would embody the generally accepted properties of this notion, and to do something on that basis”. The axioms have something of the flavor of the final definition of “polyhedron” in Lakatos [1976]. A “computation” is a function on hereditarily finite sets. Yet the axioms on computations are perfectly reasonable. Indeed, they are obvious truths about the notion of computability in question.

In sum, the notions of (idealized) human computability (and idealized machine computability) are now about as sharp as the notion of a natural number. There is not much room for open-texture anymore. The present conclusion is that this notion is the result of the last seventy years of examining texts like Turing [1936] and the overwhelming ensuing success of the theory of computability. It is not accurate to think of the historical proofs of CT and related theses as establishing something about absolutely sharp pre-theoretic notions. Rather, the analytical and mathematical work served to sharpen the very notions themselves. Sieg says as much, noting that the analyses result in a “sharpening [of] the informal notion”. As above, this is the norm in mathematics.