

What is the Church-Turing Thesis?

Udi Boker and Nachum Dershowitz

Abstract We aim to put some order to the multiple interpretations of the Church-Turing Thesis and to the different approaches taken to prove or disprove it.

[Answer:] Rosser and its inventor proved that its beta-reduction satisfies the diamond property, and Kleene (pron. clean-ee) proved that it was equivalent to his partial recursive functions. The previous result combined with a similar one with the Turing Machine, led to the Church-Turing thesis.

[Question: "What is . . . ?"]

—Quizbowl Tournament (2004)

1 Introduction

The notions of *algorithm* and *computability* as formal subjects of mathematical reasoning gained prominence in the twentieth century with the advent of symbolic logic and the discovery of inherent limitations to formal deductive reasoning. The most popular formal model of computation is due to Alan Turing. The claim that computation paradigms such as Turing's capture the intuitive notion of algorithmic computation, stated in the form of a thesis, is roughly as follows:

Church-Turing Thesis: *The results of every effective computation can be attained by some Turing machine, and vice-versa.*

Udi Boker

School of Computer Science, The Interdisciplinary Center, Herzliya, Israel, e-mail: udiboker@idc.ac.il

Nachum Dershowitz

School of Computer Science, Tel Aviv University, Ramat Aviv, Israel, e-mail: nachum@tau.ac.il

Any other standard full-fledged computational paradigm – such as the recursive functions, the lambda calculus, counter machines, the random access memory (RAM) model, or most programming languages – could be substituted here for “Turing machine”. Turing’s model has the advantage, however, of appealing only to the most basic of actions, namely, reading, writing, testing, and moving, as does Emil Post’s essentially identical and contemporaneous notion of computation [100]. (Other paradigms may model other aspects of computation better, such as recursion with recursive functions [90] or combinatory logic for composition [47].)

In what follows, we analyze different interpretations that this thesis has been given over the years and what it might mean to prove or disprove it.

Historically, the thesis was first formulated in the thirties. It started with a suggestion of Alonzo Church, based on the lambda calculus, in 1936: “The purpose of the present paper is to propose a definition of effective calculability” [30]. This was followed by the introduction of Turing machines, as a basic model of computability, in the same year [133]. Shortly thereafter, the claim that these two proposals both captured the intended notion was stated as a “thesis” by Stephen Kleene [73]. Church, Turing, and Kleene had each his own original interpretation of the claim, which were not identical. Over time, their positions in this regard developed [66, 39]. Moreover, the understanding of the thesis, as expressed in the scientific literature, has continued to evolve.

Our intention here is not to get into the historical question of how exactly Church or Turing themselves interpreted the thesis, but rather to analyze its different interpretations today. For historical aspects of the thesis, one can consult [126, 66].

The Church-Turing Thesis itself is extensional, speaking of *what* can be effectively computed, whereas the claims for and against it are intensional, arguing about *how* a computation can be accomplished. We examine first the extensional claim, looking at what type of entities are meant to be computed.

The *Extended Church-Turing Thesis* [93], otherwise known as the *Invariance Thesis* [46], is more presumptuous. It asserts, not only that there is a Turing machine corresponding to each effective algorithm, but that there is one that is no more than polynomially slower (counting atomic actions) than any other effective means of performing that computation. We do not consider this extended thesis here and concentrate on the more fundamental computability question.

2 What is Computed?

The thesis asserts that there exists some Turing machine that has the same outcomes as any particular effective calculation.

Yet, there are different views as to what it is exactly that a Turing machine computes.

What Does a Turing Machine Compute?

The input of an ordinary Turing machine is a finite tape consisting of a sequence of cells, each containing one letter taken from a finite alphabet Σ . The tape extends itself with blanks whenever the machine attempts to go beyond its current far end. The computation may also involve finitely many additional symbols, not in Σ . If and when the computation halts, the same tape contains the output – perhaps followed by trailing blanks that are not considered part of the output. By this very standard view, a Turing machine M computes a *partial function* $f_M : \Sigma^* \rightarrow \Sigma^*$ over *finite strings*, partial because some machines do not halt for some inputs. This is the version of Turing machine, out of many, that is most relevant to the intent of the original thesis.

All the same, different views are in place, even in Turing’s classic 1936 paper. There, Turing analyzed the computability of real numbers by means of his machines – using their infinite expansions to represent them:

Some of the symbols written down will form the sequence of figures which is the decimal of the real number. . . . The others are just rough notes to “assist the memory”. It will only be these rough notes which will be liable to erasure. [133]

From this point of view, a machine M runs forever, generating an infinite string. If the content of every cell of the tape eventually stabilizes, then the infinite sequence of stabilized cell values is deemed its output. Accordingly, such a machine computes a *partial function* $f_M : \Sigma^* \rightarrow \Sigma^\omega$, where Σ^ω is the set of infinite strings (padded with infinitely many blanks, when necessary).

As each step in the run of a Turing machine can alter a symbol written in some previous step, the output of an infinite run can be viewed, not as one possibly infinite string in the limit – at the “end” of the computation, but rather as a sequence of finite strings, one per step. By this view, the machine computes a *total function* $f_M : \Sigma^* \rightarrow (\Sigma^*)^\omega$.

Another issue is what the tape contents actually represent.

The tape is often imagined to be infinite, either on one end or on both. Thus, it may be thought of as an infinite string – perhaps with only finitely many positions nonblank initially. According to this view, a finite computation of a Turing machine M is a partial function $f_M : \Sigma^\omega \rightarrow \Sigma^\omega$ over *infinite strings*, while an infinite computation is a partial function $f_M : \Sigma^\omega \rightarrow \Sigma^\omega$ or a total function $f_M : \Sigma^\omega \rightarrow (\Sigma^\omega)^\omega$, depending on whether one looks only for stable outputs or at the sequence of tapes.

Aside from the basic version of a Turing machine with a single tape, many variants of Turing machines have been considered in the literature, having multiple tapes and/or multiple heads with varying levels of ability to read from and write. By some views [1], a Turing machine may be considered to compute *varying* sorts of functions over finite and infinite strings, as long as the basic model at hand preserves the atomic operations and core functionality of a Turing machine.

Nevertheless, there is a significant difference between the view that a Turing machine computes a partial function over finite strings and some of the other views that involve infinite entities. In the former, the computation producing an output is finite, whereas in the latter, the computation is expected to go on forever. Our interest here is in finite but unbounded computations alone.

What Entity Does an Effective Computation Compute?

A *quick answer* might be “a function over strings”, since a Turing machine computes with strings, and we aim at comparing something to it. However, as is well known, Turing machines are computationally equivalent to many other computational models, such as the recursive functions, the lambda calculus, and modern programming languages, all of which compute functions over different domains, such as natural numbers for the recursive functions and lambda terms for the lambda calculus. Accordingly, the thesis may – and sometimes is – stated with respect to one of these formal models of computation rather than in terms of Turing machines.

Yet, how can one compare computational models that operate over different domains? The original approach of the 1930s was to consider some “natural mappings” between these domains. A more thorough approach is taken in [13], analyzing how the choice of mapping between domains can influence the comparison of computational models. Following the latter approach to comparing computational power, it is shown in [15] that “Turing computability” can be formally extended from a partial unary function over strings to a partial function over an arbitrary structured countable domain by using any structure-preserving injective mapping between this countable domain and the natural numbers. It is guaranteed that regardless of the choice of such a mapping, it cannot influence the comparison result.

A *formal answer*, then, is that an effective computation is a partial function over an arbitrary structured countable domain, and that the thesis can be stated with respect to functions over natural numbers.

A different answer is that an effective computation can be any entity, as long as we find the computation effective, regardless of whether or not the domain is structured. Under this interpretation, if we agree on the existence of such an effective computation, we should either look into an appropriate method of comparing Turing machines to such effective computations, or, if we claim that they are incomparable, we must simply admit that the thesis is wrong. This approach can be further classified into two variants: (i) considering only functions, yet over domains that are uncountable or unstructured; and (ii) considering entities that are not functions.

Regarding a function over an uncountable or unstructured domain, one may argue whether it can at all be effective. The prevailing view is that it cannot – an unstructured domain is just an arbitrary set with no differentiation between its elements, so a function over it has little meaning. As for an uncountable domain, one should consider the question of whether the thesis involves a [human] agent (as elaborated on in the next section). If it does, the dominant view is that the agent’s ability to distinguish between different input or output elements is limited to a countable domain [38, 97].

As for a computation that is not a function, a common example in the literature is that of *interactive computation*. It has been claimed that an interactive computation, even a sequential one, does not compute a function and is therefore not Turing computable in the above sense [58, 59].¹ There is, however, a different view of

¹ Peter Wegner and Dina Goldin asserted that the Church-Turing Thesis should be interpreted with respect to functions over natural numbers [58]. Yet, if one considers it in the more general form,

interactive computations, under which they are comparable to Turing machines [1]. By this view, the claim of [58, 60] is a syntactic one, not much different from saying that the recursive functions are not Turing computable because they operate over numbers and not over strings. A sequential interactive computation is considered under this view as nothing more than a function over finite strings – an interactive computation I computes in each iteration a function $f_I : \Sigma^* \rightarrow \Sigma^*$, where the input is the string produced by the previous iteration, concatenated with a fresh environment input. The input string would also contain some information on the inner status of the computing device. As f should be the same function for all iterations, we may simply ask whether f is Turing computable (cf. [23, Defs. 14–15]). Turing already considered this setting in his classic 1936 paper, yet concentrated on “automatic machines”, as opposed to nondeterministic machines, or possibly interactive machines, which he termed “choice machines”:

For some purposes we might use machines (choice machines or c -machines) whose motion is only partially determined by the configuration. . . . When such a machine reaches one of these ambiguous configurations, it cannot go on until some arbitrary choice has been made by external operator. [133]

Turning to choice machines, a natural question is what about *random* computations that do not always return the same output for the same input, and thus do not compute a function, yet are guaranteed to return the correct answer with arbitrarily high probability. A prevailing answer is that such a computation may be regarded as computing a function, with the most likely outcome as its value. See more on randomness in Section 6.2.

The Bottom Line

The dominant view is that the core question underlying the thesis concerns functions of natural numbers, or – more precisely – functions over (string) representations of numbers. Computations of different entities are to be considered, by this view, as built over numerical encodings. This makes the extensional claim clear and simple, allowing us to concentrate on the core issue of how can an effective computation be performed. By this view, we may rephrase the Church-Turing Thesis as follows:

Church-Turing Thesis: *The partial functions over the natural numbers that can be effectively computed are exactly the functions that can be computed by Turing machines.*

they claim that he/she must admit it wrong due to the actuality of interactive computations. All this means is that the original thesis needs to be adapted to refer to the nature of non-function-computing uses of computation. See, for example, [23].

3 How is it Computed?

There are conflicting views as to what computational means can participate in an effective computation. A first question is *who/what* makes the computation: Is it a *human being* or a *machine*? The next questions correlate between the two. If a human performs the computation, what physical entities can he/she employ? And if the computation is done by a machine, must a human be able to provide the input to the function and also understand its output?

The dominant view is that a human (or some other intelligent agent) should be able to use the computation for her own purposes. A natural phenomenon that computes some function over entities that we, qua humans, cannot figure out is possibly an interesting issue to investigate in physics, but it is somewhat off the Church-Turing track [38, 97]. (It relates more to what is sometimes called *Deutsch's Principle*, which we discuss briefly in Section 6.2.2. As opposed to the *Physical Computational Thesis* (see below), which demands that a human be able to communicate and understand the input and output of a computation, this principle relates to physics alone, its capacities and limitations, without getting into what humans might be able to derive from such computations.) The main question here is whether the human performs the computation with his/her own “bare hands” or with the possible aid of some physical machinery.

The classic minimal means for a human to use in a computation is (unlimited) paper and pencil. Turing, in his classic 1936 paper, considered these quotidian means when providing arguments in favor of Turing machines.

When allowing the usage of additional machinery, there are obviously debates over what kind of machinery should be allowed. This argument is therefore related to the intensional question of how a computation is to be performed. The general interpretation permits arbitrary physical devices as long as their construction accords with modern physics. We elaborate on physical aspects in Section 6.

The next issue is how creative can the human participant be. Undoubtedly, no limits are put on the ingenuity and creativity of the human who *designs* a solution. On the other hand, when *performing* the computation, his or her actions should be systematic. The prevailing view is that, in the course of performing the computation, the human should follow definite, unambiguous, predefined instructions, and avoid any appeal to creativity [76]. When considering the question of whether the human mind is limited to Turing computability, the creative aspect of human activity also comes into play, a subject we will touch on later.

The Bottom Line

Both interpretations, that of only having paper and pencil and that of having arbitrary machinery at one's disposal, are widely considered in the literature. Both are referred to as the “Church-Turing Thesis” and sometimes under other names, too. We shall

distinguish between the *mathematical* and *physical* theses, respectively.² Although we use the adjectives “mathematical” and “physical”, the theses do not speak of the limitations of mathematics and physics per se, but rather of the capabilities of humans who use those resources for performing computations.

Mathematical Computational Thesis: *The partial functions over the natural numbers whose computations can be effectively performed by a human who systematically uses only paper and pencil of finite but unlimited quantity and who has finite but unlimited time are exactly the Turing-computable numerical functions.*

Physical Computational Thesis: *The partial functions over the natural numbers whose computations can be effectively performed by a human who systematically uses any physical device of finite but unlimited resources, including finite but unlimited time, are exactly the Turing-computable numerical functions.*

We are taking a purely anthropocentric point of view: What can a human compute systematically, following predefined rules, be it with only the most basic of tools (pencil and paper), or with arbitrarily elaborate physically feasible devices? What it may be that natural systems or fabricated apparatuses can calculate for their own, internal consumption is not the issue under inquiry, to the extent that they do not contribute to repertoire of functions available to us humans.

We elaborate on the mathematical thesis in Section 5 and on the physical thesis in Section 6.

4 What Can Be Proved?

To Prove or Not to Prove? This is the Philosophical Question

There is a contentious philosophical debate on whether there is a sense in trying to prove the thesis. (See, for example, the arguments in [115, 50] and the references therein.) In a nutshell, the dominant claims against a proof assert that an intuitive notion, such as effectiveness, cannot be proved equal to a formal notion, such as Turing computability. No matter what proof you provide, it is based on some basic assumptions. Therefore, the “proof” only transforms the Church-Turing Thesis to a thesis claiming that one’s assumptions properly capture the essence of effective computation. On the other hand, claims in favor of provability assert that every

² There are numerous other names for the various versions of the thesis in the literature. Even the terms “mathematical” and “physical” have different connotations for different authors.

proof takes some “first principles” as axioms, or else nothing is provable. Hence, agreeing on first principles for effective computability, on top of which a proof of the thesis is considered, is as reasonable as agreeing on first principles to logic and mathematics, on top of which every mathematical proof is considered. One can still argue about how close the principles are to the intended intuitive notions, just as one may speculate whether the axioms of ZFC properly capture the full essence of a set.

In parallel with the philosophical debate on whether or not a proof of the thesis is plausible, there is a rich literature on *arguments* for and against the thesis, as well as various *concrete proofs* of the thesis within specific formal settings. We elaborate on some of them in Sections 5 and 6.

4.1 Against Provability

The view that the thesis cannot be proved harks back to Church and Kleene:

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 notion. [30, p. 356]

Since our original notion of effective calculability of a function (or of effective decidability of a predicate) is somewhat vague intuitive one, the thesis cannot be proved”. [74, p. 317]

Likewise, Janet Folina claims, “There is a good deal of convincing evidence that CT [the Church-Turing thesis] is true”, yet, “It is simply not possible to prove that an informal, intuitive, notion has a single precise (formal or axiomatic) articulation” [50, p. 321].

László Kalmár further claims that, not only doesn’t the thesis properly capture the intuitive notion of effectiveness, but we must not try to mathematically capture it:

There are pre-mathematical concepts which must remain pre-mathematical ones, for they cannot permit any restriction imposed by an exact mathematical definition. Among these belong, I am convinced, such concepts as that of effective calculability, or of solvability, or of provability by arbitrary correct means, the extension of which cannot cease to change during the development of Mathematics. [70, p. 79]

As for Turing, he also considered computability to be an intuitive notion that cannot be defined mathematically, though he did claim that the right viewpoint to the question is via analyzing the possible processes of a computation:

No attempt has yet been made to show that the “computable” numbers include all numbers which would naturally be regarded as computable. All arguments which can be given are bound to be, fundamentally, appeals to intuition, and for this reason rather unsatisfactory mathematically. The real question at issue is “What are the possible processes which can be carried out in computing a number?” [133, p. 249]

4.2 In Favor of Provability

Kurt Gödel has been reported (by Church in a letter to Kleene cited by Davis in [37]) to have thought that

It might be possible . . . to state a set of axioms which would embody the generally accepted properties of [effective calculability], and to do something on that basis.

Joe Shoenfield similarly asserted that:

It may seem that it is impossible to give a proof of Church's Thesis. However, this is not necessarily the case. . . . In other words, we can write down some axioms about computable functions which most people would agree are evidently true. It might be possible to prove Church's Thesis from such axioms. [118, p. 26]

The late Elliott Mendelson further claimed:

It is completely unwarranted to say that CT is unprovable just because it states an equivalence between a vague, imprecise notion (effective computable function) and a precise mathematical notion (partial recursive function). [87, p. 232]

He asserts that a nonmathematical notion can indeed be fully matched with a mathematical one, as for example is the case with what he calls *Peano's thesis* [87, p. 230], which provides an exact mathematical definition in terms of sets of ordered pairs for the intuitive notion of a function.

Mendelson (and Gandy [54] before him) even asserts that Turing's arguments in favor of the thesis (see Section 5) are in fact a proof for it:

The fact that it is not a proof in ZF or some other axiomatic system is no drawback; it just shows that there is more to mathematics than appears in ZF. [87, p. 233]

Shapiro supports Mendelson's view, claiming, "We can and do prove theses a lot like CT" [115]. (For a lengthy discussion of the issues, see [116].)

Mendelson further claims that the concepts of effectiveness and Turing computability (recursiveness) are equally vague:

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.

He compares it to the case of functions and sets:

Functions are defined in terms of sets, but the concept of set is no clearer than that of function.

In 2000, Harvey Friedman predicted that

There will be [in the 21st century] an unexpected striking discovery that any model of computation satisfying certain remarkably weak conditions must stay within the recursive sets and functions, thus providing a dramatic "proof" of Church's Thesis. [51]

4.3 What Does it Mean to Disprove the Thesis?

As with the positive direction of proving the thesis, there are also philosophical debates on whether there is a sense in trying to disprove it. If we deny the ability to formalize an intuitive notion, neither can we prove that something is included in it.

Yet, the possibly more common view is that there is a significant difference between the essence of proving and disproving the thesis. Indeed, once an effective computation scheme that is stronger than Turing machines is introduced, we might agree on its effectiveness and recognize the falsehood of the thesis.

Saul Kripke, for example, writes,

That the recursive or Turing-computable functions are all in fact effectively calculable, can I think easily be established simply on the basis of an intuitive notion of calculability. No axiomatization is required, just a direct argument. And the argument should be regarded as a rigorous intuitive proof of its result. [81, p. 78]

It is harder, no doubt – some would say impossibly harder, to prove that *all* effective means are included in a proposed formalism than that some particular formally defined function is effectively computable.

Primitive recursion, for instance, was considered at some point as a candidate formalism for characterizing effective computability, yet once general recursion was introduced, it was widely agreed to be effective and to exceed the computational power of primitive recursion.³ Quoting Henk Barendregt [8]:

One may wonder why doubting Church's Thesis is not a completely academic question. This becomes clear by realizing that [Skolem in 1923] had introduced the class of primitive recursive functions that for some time was thought to coincide with that of the intuitively computable ones. But then [Ackermann in 1928] showed that there is a function that is intuitively computable but not primitive recursive.

5 The Mathematical Thesis

What are the "Rules of the Game"?

The exact rules of the game are, of course, vague, as is the abstract notion of effectiveness. Still, some rules are clear enough. Computation is discrete, that is, done in steps, and the number of steps should be finite, though it may be unbounded. The computational process has a finite description, independent of the specific inputs. The space and time made available for the purpose of the computation are accordingly unbounded but finite. Computation is generally considered to be symbolic. "There is no calculation without representation" [49] (see [29]). That is, the data with which one works consist of finite symbolic configurations, where the symbols (or labels) are

³ In [13, 16] it is shown that while strict containment of function sets is in general not enough to infer the presence of extra computational power, general recursion is indeed more powerful than primitive recursion, even under rigorous definitions of power comparison.

drawn from some finite set given in advance. It bears stressing that computation of a numerical or string function may involve additional domains, which need, likewise, to be represented symbolically.

This ideal, unbounded, nature of the computation – of both a Turing machine and an “effective computation” – is intuitively in tension with the effectivity of the computation. Indeed, in contrast to the common claim that Turing-computability is effective, there are counterclaims that even Turing machines are not effective, mentioned below. Yet, “effective” was never meant in the sense of reasonable or practical demands on time, or space, or other resources. This is the reason for the rich field of complexity theory, and to the distinction between computability and complexity. The latter relates to the extended Church-Turing thesis, while we concentrate on the former.

The Empirical Evidence

Since the seminal works of Church, Turing, Gödel, Herbrand, Kleene, and Post in the thirties, defining general computational models, and until current days, all of the general-purpose models were shown to have exactly the same computational power given finite and unbounded time and memory. Among these models are Turing machines, the recursive functions, lambda-calculus, Post canonical systems, modern programming languages, and many others. This empirical fact, having so many different views of computation, all leading to the same set of functions, gave significant confidence in the validity of the Church-Turing thesis. As Hartley Rogers writes in his classic text in support of the thesis:

The proposed characterizations of Turing and of Kleene, as well as those of Church, Post, Markov, and certain others, were all shown to be equivalent. [108, pp. 18–19]

5.1 Non-empirical Arguments in Favor of the Thesis

Next, we enumerate some of the efforts made to justify or prove the thesis on a non-empirical basis, starting with Church’s and Turing’s original ones.

Church’s Arguments

Church, in his classic paper, provided two arguments in favor of identifying effectiveness with λ -definability or general recursiveness. The first concerns an effective algorithm and the second an effective logical proof:

No more general definition of effective calculability than that proposed above can be obtained by either of two methods which naturally suggest themselves (1) by defining a function to be effectively calculable if there exists an algorithm for the calculation of its values (2) by defining a function F (of one positive integer) to be effectively calculable if, for every positive integer m , there exists a positive integer n such that $F(m) = n$ is a provable theorem. [30, p. 358]

In the same paper, it is shown that λ -definability and general recursiveness are equivalent in the sense that each can simulate computation in the other sense.

In the first argument for comprehensiveness of the definition of effectiveness, Church considers the following:

- An algorithm consists of a finite sequence of expressions: “For example, in the case of a function F of one positive integer, an algorithm consists in a method by which, given any positive integer n , a sequence of expressions (in some notation) $E_{n1}, E_{n2}, \dots, E_{nr_n}$, can be obtained”.
- Each step of the algorithm should be recursive: “We take the effective calculability of G and H to mean recursiveness” [30, p. 357], where G and H are functions over natural numbers that represent the algorithm’s step-computation when representing the algorithm’s expressions with their Gödel numbers.

Under these assumptions, “the recursiveness (λ -definability) of F follows by a straightforward argument” [30, p. 357]. Church further claims:

If this interpretation or some similar one is not allowed, it is difficult to see how the notion of an algorithm can be given any exact meaning at all. [30, p. 357]

Overall, the logic community was not at all convinced by Church’s arguments, as demonstrated by Gödel’s response (reported by Kleene):

According to a November 29, 1935, letter from Church to me, “Gödel regarded as thoroughly unsatisfactory” Church’s proposal to use λ -definability as a definition of effective calculability. [75]

In the second argument, Church considers a proof to be in some “symbolic logic, which contains a symbol . . . for equality”, starting with “*formal axioms*” and proceeding by “the *rules of procedure*”. Analogously to the first argument, he assumes that “each rule of procedure must be a recursive operation”, and that “the complete set of formal axioms must be recursively enumerable” [30, p. 357]. Under these assumptions, the recursiveness of the entire proof directly follows.

The main claim against this argument is that it’s circular, demanding that the basic steps of the computation be themselves recursive (e.g. [118, 120, §6.5]). All the same, similar arguments were given by Turing (in addition to other arguments), and they have been further developed by Kripke [81]. See below.

Turing’s Arguments

In his classic paper, Turing gave three different arguments in favor of identifying effective computability with computability via Turing machines. The first builds on the limitations of a human being working in a rigorous, bureaucratic way. His second argument concerns the equivalence of Turing computability and “Hilbert functional calculus”, and resembles Church’s arguments, namely, “The steps of any effective procedure (governing proofs in a system of symbolic logic) must be recursive” [120]. This has been further developed by Kripke to reduce the mathematical thesis to a thesis about first-order logic (see below). The third provides “empirical evidence” that large classes of real numbers are Turing computable.

We elaborate on Turing's first argument. It builds on the following limitations of a [human] computer [133, pp. 249–250]:

- “Computing is normally done by writing certain symbols on paper.”
- “The two-dimensional character of paper is no essential of computation. I assume then that the computation is carried out on one-dimensional paper, i.e. on a tape divided into squares.”
- “The number of symbols which may be printed is finite. If we were to allow an infinity of symbols, then there would be symbols differing to an arbitrarily small extent.”
- “The behaviour of the computer at any moment is determined by the symbols which he is observing, and his ‘state of mind’ at that moment.”
- “There is a bound B to the number of symbols or squares which the computer can observe at one moment.”
- “The number of states of mind which need be taken into account is finite. The reasons for this are of the same character as those which restrict the number of symbols. If we admitted an infinity of states of mind, some of them will be ‘arbitrarily close’ and will be confused.”
- “In a simple operation not more than one symbol is altered. . . . The squares whose symbols are changed are always ‘observed’ squares.”
- “The new observed squares must be immediately recognisable by the computer. . . . They can only be squares whose distance from the closest of the immediately previously observed squares does not exceed a certain fixed amount.”

Turing's argument is considered by many to be most convincing one, some even going so far as to view it as an actual proof. For example, Robin Gandy called it “Turing's proof”, and said, “The proof is quite as rigorous as many accepted mathematical proofs . . .” [54, p. 82], and likewise Mendelson [87] and Shapiro [115].

Church, reviewing Turing's classic paper, was convinced by the arguments: “Computability by a Turing machine . . . has the advantage of making the identification with effectiveness in the ordinary (not explicitly defined) sense evident immediately – i.e. without the necessity of proving preliminary theorems” [31, p. 43].

Gödel adopted Turing's argument, saying [136, p. 203], “The resulting definition of the concept of mechanical by the sharp concept of ‘performable by a Turing machine’ is both correct and unique.” In later years, he did, however, raise doubts about Turing's arguments:

What Turing disregards completely is the fact that mind, in its use, is not static, but constantly developing. . . . Although at each stage the number and precision of the abstract terms at our disposal may be finite, both (and, therefore, also Turing's number of distinguishable states of mind) may converge toward infinity in the course of the application of the procedure. [56, p. 306]

The common response to this objection of Gödel is that an algorithm should be completely given in advance. Kleene wrote:

Thus algorithms have been procedures that mathematicians can describe completely to one another in advance of their application for various choices of the arguments. How could

someone describe completely to me in a finite interview a process for finding the values of a number-theoretic function, the execution of which process for various arguments would be keyed to more than the finite subset of our mental states that would have developed by the end of the interview, though the total number of our mental states might converge to infinity if we were immortal. [76, p. 50]

Wang notes that

Gödel realizes that his incompleteness theorem does not by itself imply that the human mind surpasses all machines. [135]

There is no reason why the number of states of the mind should not converge to infinity in the course of its development. [134, pp. 325–326]

In other words, human mathematical reasoning – as opposed to algorithmic reasoning, by human or machine – need not be held to the requirement that states of mind be bounded [112]. The assertion that humans or humankind *cannot* outperform Turing machines even in such non-algorithmic ways has been referred to as the “Non-mathematical Church-Turing Thesis” [7], to avoid confusion with the algorithmic, mathematical one we are discussing.

Another criticism of Turing’s arguments, or more accurately a sort of request for further development, is that they are not formal enough. The quest is for a set of postulates that define the limitations of effective computability along with a proof that under these limitations no computation exceeds Turing computability. There are several works along this line, and we elaborate on some of them next.

Kripke’s Arguments

Saul Kripke reduces the Church-Turing Thesis to “Hilbert’s Thesis” (as Martin Davis called it [9, p. 41]), namely, “The steps of any mathematical argument can be given in a language based on first-order logic (with identity)” [81, p. 81]. He asserts, as did others before him,⁴ that all “computation is a special form of a mathematical argument”, and that “computation is a deductive argument from a finite number of instructions. . . albeit one of a very specialized form” [81, p. 80].

Thus, by Hilbert’s Thesis, the premises of the computation as well as its steps can be expressed in first-order logic. This assumption in hand, Kripke obtains, as a corollary of Gödel’s completeness theorem for first-order logic (with identity), the *First-Order Algorithm Theorem*, as he calls it:

Any algorithm whose steps are meaningfully expressed in a first-order language, the conclusion of the computation does indeed follow from the premises. [80, at 1:18]

Since every first-order proof is Turing computable, as already shown by Turing, a proof of the thesis directly follows. “What [Turing] gives as argument 2 in his paper is essentially that any first-order algorithm can be computed by a Turing machine, as expressed in my own terminology” [80, at 1:19]. The first-order inference system works with terms and formulas to derive the output of the algorithm.

⁴ Sieg [119] in explaining Church [30]: “Calculability is explicated by that of derivability in a logic.”

It should be remarked that Kripke further claims that the proof also holds with respect to machine computations, as their computation steps must follow some deductive rules:

Any machine – think of it as an abstract object, it is a program. Do the steps in the program follow from the initial instructions in the program as given or don't they? And can they be expressed in a final conventional mathematical language formalized in first order logic or can't they? If the answer is no, I don't regard it as a computation at all. If the answer is yes, the theorem that I just stated is enough to give an affirmative answer that the function that comes out of it must be Turing computable. [80, at 1:27]

An objection to Kripke's argument, for example by Michael Rabin [80, at 1:54], is that it is circular: A computation can be viewed as a logical proof, and the proof should be susceptible to mechanical verification. As Kripke puts it: "Granted that the proof relation of such a [formal logical] system is recursive (computable)" [81, p. 81]. Yet, disbelieving the thesis, there might be a stronger logic that can be mechanically checked (by an effective machine stronger than a Turing machine). Rabin would be equally surprised if a machine stronger than Turing's is found or if a notion of computation not expressible in first-order logic is found.

Sieg's Arguments

Gandy [53], Wilfried Sieg [121], and Robert Soare [126] divide Turing's first argument into a philosophical part, in which there are assumptions on what an idealized human (a "computer" with an *o*) can rigorously do with paper and pencil, and a mathematical part, in which it is proved that under the above philosophical assumptions, every computation is Turing computable.

The assumptions on the computer are reduced by Sieg to principles of boundedness and locality:

The constraints Turing imposes on symbolic processes derive from his central goal of isolating the most basic steps of computations, that is, steps that need not be further subdivided. This objective leads to the normative demand that the configurations, which are directly operated on, must be *immediately recognizable* by the computer. This demand and the evident limitation of the computer's sensory apparatus motivate most convincingly two central restrictive conditions.

(Boundedness) A computer can immediately recognize only a bounded number of configurations.

(Locality) A computer can change only immediately recognizable configurations. [121]

Sieg's formal axioms aim at capturing these informal assumptions.

The axioms are formulated for discrete dynamical systems and capture the above general ideas precisely; they should be viewed as determining classes of "algebraic structures" of which particular models of computations are instantiations. [121]

Sieg considers a computation to be a state transition system, in which states (as in [53]) are finite objects that can be represented by nonempty hereditarily finite sets over an infinite set of atoms. The classes of states that are considered are closed

under isomorphism. His axioms formalize the locality and boundedness principles in this setting.

If x is a given state, regions of the next state are determined *locally* from particular *parts* for x on which the computer can operate. *Boundedness* requires that there are only finitely many different kinds of such parts. [121]

These “kinds” are formalized by isomorphism types, termed *stereotypes*.

A Turing computer is defined to comprise a class S of states closed under isomorphism, a finite set T of stereotypes, and an operation G on the stereotypes of T , satisfying the requirement that for every state x of S there exists a state z of S (the “next state”) such that the following three axioms hold:

- ◇ There exists in x a unique maximal stereotype from T . This unique stereotype is the *causal neighborhood* of x .
- ◇ Applying the operation G on the causal neighborhood of x yields, up to isomorphism, a unique region of the state z .
- ◇ The state z consists of its above region and the whole of x , excluding the causal neighborhood of x .

The first two axioms formalize the principal of *local causality* and the third the principle of *global assembly*.

A computation of a Turing computer is defined to be a finite sequence of transitions between states using operation G , *halting* when the next state is the same as the preceding one.

As for computations over domains that do not consist of hereditarily finite sets, Sieg speaks of “suitable encoding and decoding” without detailing what is to be considered suitable:

A function \mathbf{F} is (Turing) *computable* if and only if there is a Turing computer \mathbf{M} from whom computation results one can determine – under suitable encoding and decoding – the values of \mathbf{F} for any of its arguments. [121]

Sieg’s work served to clarify and sharpen Turing’s arguments. Some researchers view the Turing-Sieg arguments as actually constituting a proof of the thesis. Sieg, himself, views his work as proving a representation theorem for the class of systems that can be formulated in a way that meets the stated criteria. Unsurprisingly, those who are not convinced by Turing’s arguments tend not to be convinced by Sieg’s more formal arguments either. Another criticism of Sieg’s approach (and of Gandy before him) is that expressing computability axioms in terms of hereditarily finite sets representing states – however sensible under the standard set-theoretical, foundational view of the mathematical universe⁵ – is nonetheless insufficiently general for faithfully modeling effective computations [63, 81].

⁵ Gandy’s idea of basing a formalization of computability on hereditarily finite sets, which are precisely the finite objects of set theory, was presaged by Moto-o Takahashi [131].

Gurevich, Boker, and Dershowitz's Arguments

As Church argued in his classic paper [30], a possible take on the mathematical thesis is that a Turing machine can compute every function that can be computed algorithmically, where the term “algorithm” stands for the intensional means of a computation. A central point of note is that the equivalence claim is extensional and not intensional – the algorithm itself is not supposed to be captured by the computation of a Turing machine and may have various abstraction levels. It is only the result of the algorithm, namely the function it computes, that is supposed to be captured by the result of the Turing machine.

When aiming at proving the thesis, one approach is to first define what an algorithm is [90, 63], and what atomic operations it can effectively employ.

At the turn of the millennium, Yuri Gurevich proposed three basic postulates for sequential (non-parallel and non-interactive) algorithms over arbitrary mathematical structures [62]. By this view, every sequential algorithm should satisfy the following properties:

- ◊ It can be viewed as a set of states and a transition function from state to state. The importance of this generic view has been emphasized by Don Knuth [77].
- ◊ Its states are (first-order) structures sharing the same finite vocabulary. States are closed under isomorphism, and the transition function preserves isomorphism. The relevance of logical structures appears already in Emil Post’s notes [101]; the importance of isomorphism is stressed by Gandy [53].
- ◊ There is a finite bound on the number of vocabulary-terms that affect the transition function. This formalizes ideas expressed by Kolmogorov [78].

Based on these three principles, Gurevich proved that every such sequential algorithm is captured by an *abstract state machine*, a formalism he defined that expresses generic algorithms as sets of parallel conditional assignments employing the operations of the structure [62].

Gurevich’s postulates and proof are a natural starting point for axiomatizing effectiveness over arbitrary domains, not just strings. Yet, there are a couple of “loose edges” that allow an algorithm that satisfies Gurevich’s postulates to yield an ineffective computation:

- The atomic operations, which can be viewed as the initial state of the algorithm, are not limited and might thus introduce ineffectiveness.
- The postulates consider a single algorithm and not an entire computational model. Interpreted over an arbitrary domain, the computed function might be ineffective with respect to a different domain or even with respect to the same domain, under a different representation.

In 2005, Boker and Dershowitz addressed the above issues by considering the whole of a *computational model*, comprising the computations of a *set* of functions, rather than a single function, requiring all computations within the same model to share the same domain view, and restricting the initial states to only have finite data. The additional restrictions were set out in a fourth postulate [12]:

- ◊ There is an initial state to the computation and it comprises only finite data in addition to the domain representation. The latter is isomorphic to a Herbrand universe.

Boker and Dershowitz proved that every computational model that satisfies the above four postulates cannot compute more than Turing machines [15].

In 2008, Dershowitz and Gurevich addressed the two issues above in a slightly different fashion [43]. They required an injective mapping between the arbitrary domain and the natural numbers, and limited the initial data to be tracked, under that representation, by a recursive function.

Shortly afterwards, Boker and Dershowitz showed that the three different approaches taken to define effectiveness on top of Gurevich's postulates for a sequential algorithm – their's from 2005, Dershowitz and Gurevich's from 2008, and Wolfgang Reisig's from 2008 [105], yield the exact same notion of effectiveness [17]. Whereas the latter two involve a measure of circularity, using recursiveness or Turing-computability to guarantee that initial states are effective (see the objection in [122, Postscriptum]), the former, equivalent approach builds only on first principles.

In favor of their view is their notion of state, namely, any mathematical first-order structure, which might be considered more general and natural than strings or hereditary finite sets. Furthermore, their axioms are built on top of Gurevich's postulates for sequential algorithms, which aim at capturing the essence of classical algorithms. As a result, their formalization of effective computability captures a most general setting of computational models. (Supportive reactions include [88, p. 351n] and [81].)

The natural criticism of this attempt is that the result just “moves the goalpost”, shifting the question to whether the provided postulates capture all of what is effective (e.g. [104, p. 431], [122, Postscriptum]). Gurevich himself believes that his postulates properly capture the notion of a sequential algorithm, but avers that no definition can capture what a more “general” algorithm is [63]. Another criticism is that the postulates are not sufficiently formal [122, Postscriptum], but that can be rectified.

5.2 Disproving the Thesis

One way to disprove the mathematical thesis would be to find a function that is not Turing-computable but which one can convincingly claim can be computed effectively by an idealized human (immortal, untiring, unerring, etc.) equipped with pencil and paper and following predetermined instructions. This would be analogous to the rejection of primitive recursion as capturing all effective computation on the basis of the effectivity of Ackermann's function.

As popular textbooks explain:

If one ever found a procedure that fitted the intuitive notions, but could not be described by means of a Turing machine, it would indeed be of an unusual nature since it could not possibly be programmed for any existing computer. [69, p. 80]

It is theoretically possible, however, that Church's thesis could be overthrown at some future date, if someone were to propose an alternative model of computation that was publicly acceptable as fulfilling the requirement of finite labor at each step and yet was provably capable of carrying out computations that cannot be carried out by any Turing machine. No one considers this likely. [83, pp. 168–169]

There have been several attempts in this direction, as we will see in this section. One of the earliest was by Kalmár [70] (explicated in [130]). In particular, he questioned the commonplace requirement that an algorithmic method be uniform for all inputs:

We regard as effectively calculable any arithmetical function, the value of which can be effectively calculated for any given arguments in a finite number of steps, irrespective how these steps are and how they depend on the arguments for which the function value is to be calculated. In particular, I do not suppose the calculation method to be “uniform”. [70, p. 73]

We will encounter such misgivings again.

Considering the computation of specific instances of the halting problem, Turing commented already in his 1936 paper as follows:

It is an immediate consequence of the theorem of §8 that δ is not computable. It is (so far as we know at present) possible that any assigned number of figures of δ can be calculated, but not by a uniform process. When sufficiently many figures of δ have been calculated, an essentially new method is necessary in order to obtain more figures. [133, p. 253]

In other words, the thesis places no limits on non-uniform human problem solving abilities, let alone those of endless generations of humankind. Alfred Tarski is quoted by Benjamin Wells [137] as referring to problems for which there is a series of non-uniform solutions as (intuitively) “decidable”.

A related debate surrounds the so called “Artificial Intelligence Thesis” and the question whether a computational device can act with general human-level intelligence. This AI thesis demands more of the mechanical device in that it must also be creative, imaginative, and intuitive like human beings. We do not address the AI question but rather concentrate on mathematical proficiency and ignore other aspects of human intelligence.

Early on, several objections were voiced that the thesis was wrong for the opposite reason, that Turing's formalism was overly rich, and that not every Turing-computable function is effective. See [96, 99]. With the growth of programming and programming languages, this view is no longer tenable and were dispensed with already in [86, 89].

Nowadays, one occasionally hears objections from the point of view that unbounded time and space are wholly unrealistic idealizations, so these imaginary machines are not “really” effective in any practical sense. From the point of view of this survey, these are tangential issues.

Another issue, first raised in [114], is how a human interprets the strings that a Turing machine uses for input and output. Michael Rescorla, for instance, has pointed out correctly that if one imposes “a deviant interpretation upon numerals then Turing machines can ‘compute’ intuitively non-computable numerical functions” [107]. See [106, 33, 107, 139] for a lively discussion. Our preferred solution is to require

that all recursive functions remain computable under whatever representation is chosen, so that any model that ostensibly computes something that a Turing machine cannot must also be capable of computing whatever a Turing machine can with the identical encoding scheme [13, 14].

Penrose's Arguments

In *The Emperor's New Mind* [94] and especially in *Shadows of the Mind* [95], the eminent physicist, Roger Penrose, contends that human reasoning cannot be captured by a mechanical device because humans detect nontermination of programs in cases where digital machines cannot [128]. Penrose thus adapts the similar argumentation of Lucas [84], which was based on Gödel's incompleteness results.

How this impinges on the Church-Turing Thesis depends on whether the means by which Penrose contends that humans can solve the specific halting problem he has in mind should be deemed "effective" or just "intuitive". This is somewhat unclear, because Penrose does not claim to know himself how it is that people achieve what he asserts they do. Were it by somehow engaging a quantum mechanism embedded in the human brain (as Penrose speculates in [95]), then this would be an argument against the physical thesis, but not the mathematical one. But if the basis of the super-algorithm used for this purpose is ordinary human mathematical knowledge, such as that taught in an undergraduate computability course (as actually employed in the diagonalization argument in [95]), then this would be a frontal attack on the mathematical thesis.

In a nutshell, Penrose's argument runs as follows:

1. Consider all current sound human knowledge about nontermination.
2. Suppose now that one could reduce said knowledge to a computer program.
3. Then one could also create a self-referential version of said program.
4. From the assumed existence of such a "diagonal" program, a contradiction to its correct performance can be derived.

Penrose's resolution of this contradiction is to deny the validity of the second step: No finite program can incorporate all the algorithmic and mathematical knowledge that finitely many humans currently have at their disposal.

Since Penrose outlines the procedure by which a human can leverage all the knowledge of the first step to determine that the program in question is nonterminating, we may conclude that his argument serves to indicate that humans are somehow more capable than Turing machines, and that the Church-Turing thesis is false.

Penrose's conclusions, namely that human reasoning cannot be faithfully simulated by Turing machines, have been heavily critiqued (e.g. [5, 28, 82]). The possible rejoinders are classified in [42]. Assuming that the *idealized* humans in question will not die before they decide what to respond and that they would neither knowingly lie nor impetuously respond without thinking (something that real humans are apt to do under the circumstances), resolutions of the conundrum raised by Penrose fall mainly in two categories: (a) humans can err in judgement while believing they are being truthful; and (b) a human may realize or suspect that there are matters that may

lie beyond one's ability to fully comprehend or to express, such as the soundness or consistency of one's own vastly complicated internal workings.

An example of objection (a) is this:

No human mathematician can claim infallibility. We all make mistakes! So there is nothing in Gödel's theorem to preclude the mathematical powers of a human mind being equivalent to an algorithmic process that produces false as well as true statements. [40]

An example of the second objection is

The possibility exists that each of the rules that a human mathematician explicitly relies on, or can be rationally persuaded to rely on, can be known to be sound and that the program generates all and only these rules but that the program itself cannot be rendered sufficiently "perspicuous" for us to know that that is what it does. . . . A program which simulated the brain of an idealized mathematician might well consist of hundreds of thousands (or millions or billions) of lines of code. Imagine it given in the form of a volume the size of the New York telephone book. Then we might not be able to appreciate it in a perfectly conscious way, in the sense of understanding it or of being able to say whether it is plausible or implausible that it should output correct mathematical proofs and only correct mathematical proofs. [103]

In other words, one may in fact only know correct rules, but at the same time be unable to establish that fact. Or, one may even realize that declaring soundness out loud might ipso facto make it false. In such cases, it would seem to us that a cautious human may simply prefer, or feel compelled, to maintain judicious silence [42].

6 The Physical Thesis

What are the "Rules of the Game"?

In the mathematical thesis, the number of computation steps is finite but unbounded. This no longer stands as a clear rule in the physical thesis – a step is not always a physical notion. We might want to consider analog phenomena and how discrete entities are conceived in the physical world. The input and output of such a computation are to be understood by the human user as natural numbers, whereas the computation and entities involved are continuous, analog processes.

We do not address the issue of what constitutes an effective real-valued function, normally expressed in terms of approximations of increasing accuracy (e.g. [117]).

Analyzing the physical thesis, one needs to look into current physical theories of the universe and how humans can accordingly take advantage of physical phenomena for performing computations. Whereas the resources consumed in a computation must be finite, though unlimited upfront, there might indeed be limitations on them stemming from the constraints of physical theories, such as the speed of light, the uncertainty principle, etc. As put by Oron Shagrir [111, p. 231]:

In this context, a physical system is any system, real or potential, such that (a) its states occupy a finite space, (b) its terminating dynamics are completed in finite real time, and (c) its dynamics are consistent with the laws of physics.

6.1 Proving the Thesis

Several arguments have been aired in support of the physical thesis. These arguments provide axioms that are claimed to express limitations on physical computations in consonance with current physical theories. Then, properties of all systems obeying those axioms are derived. In considering an attempt to disprove the physical thesis, one may ask which of the axioms is violated and argue whether a particular axiom should be adopted.

Gandy's Arguments

The most notable effort to prove the physical thesis was initiated by Robin Gandy (who was Turing's student) in 1980 [53]. He defined four principles that, according to his understanding of physics, every "discrete mechanical device" should satisfy. The principles are derived from two basic physical assumptions: that there is a lower bound on the size of atoms and that there is an upper bound on the speed of signal propagation. The main difference between Gandy's principles and Turing's assumptions on effectiveness is that Gandy allows for unbounded parallelism: There can be unboundedly many different units computing in parallel. Gandy then proved that whatever is computed by such a device is also Turing computable.

Gandy's four principles are the following, stated roughly:

- ◇ The machine can be described by a transition system, where the states are hereditarily finite sets.
- ◇ The set-theoretic rank of the states is bounded.
- ◇ There is a finite set of basic parts from which all states can be assembled.
- ◇ For each region of the machine, its next state only depends on its neighbourhood of a bounded size.

Gandy's work was much appreciated for its novelty and depth, though perceived as somewhat overly complicated [115, 63]. Sieg and John Byrnes simplified Gandy's principles and arguments [123], and they were later further modified by Sieg [121]. (See below.) They are critiqued in [111].

One natural claim against Gandy's arguments are that his four principles are too restrictive, not covering all of what is physically effective. In particular,

- The restriction to discrete devices is improper as physical devices may be analog; we should also consider what the limits of analog computations are [35, 81, 18].
- The first principle of finite information for each state and the fourth principle of local causality are not obeyed by quantum physics [72, 3].
- The parallelism is only synchronous, namely the different units proceed in synchronous steps, whereas actual independent computations are not synchronous [63].
- The requirement to have finitely many steps in a terminating computation is disobeyed by devices that allow for unbounded acceleration, which might be possible by laws of special and general relativity [35].

Sieg's Arguments

Following in Gandy's footsteps, Sieg extends the suppositions about what an idealized human "computer" can rigorously do with paper and pencil (Section 5) to assertions regarding what a parallel computing device (Gandy machine) can do [121]. These assertions still stem from the same principles of boundedness and locality, justified in this case by the physical lower bound on the size of atoms (or fundamental particles) and the physical upper bound on the speed of signal propagation, yet generalized to allow for parallelism. After formalizing these assertions, he proves that every computation of such a Gandy machine is in fact Turing computable.

Recall that Sieg defines a computer to comprise a class S of states closed under isomorphism, a finite set T of stereotypes, and an operation G on the stereotypes of T , satisfying some requirements. Now, he defines a Gandy machine to have in addition another set T' of stereotypes and another operation G' on stereotypes T' , which allow it to coordinate parallel computations. He then generalizes the axioms that formalize *local causality* and *global assembly* for such a parallel machine.

Compared to Gandy's original four principles, as roughly expressed above, Sieg retains the first and fourth, while relaxing the second and third. Since Sieg's formulation puts limits only on causal neighborhoods, entire states need not have a bounded set-theoretic rank and need not be assembled from a finite set of basic parts. Accordingly, objections to Gandy's arguments that concern his second and third principles do not hold for Sieg's arguments, while claims against Gandy's first and fourth principle, as described above, still carry over to Sieg's.

Arrighi and Dowek's Arguments

Pablo Arrighi and Gilles Dowek generalized Gandy's principle and proof to quantum computations [3].

Gandy's first principle, according to which states can be defined by finite sets, is based on his assumption that information has a finite density, which no longer holds in quantum theory. (See [4].) Arrighi and Dowek provide a different assumption that also holds in quantum theory and allows for the first principle – that each projective measurement of the system, at any given point in time, may only yield a finite number of possible outcomes.

Gandy's fourth principle of local causality is based on bounded velocity, which fails to stand in the quantum setting. Arrighi and Dowek provide instead a quantum version of causality, based on the assumption that the global evaluation of the system is "localizable", namely that it is implementable by local mechanisms [10, 110].

6.2 Disproving the Thesis

There are two types of claims against the thesis; the first is that there are currently computing devices that exceed Turing computability, and the second is that our physical theories allow for theoretical devices that exceed Turing computability and

might possibly be built in the future. The first claim, concerning existing machines, is mostly based on randomness and learnability. We have already encountered several claims that the human mind may somehow enjoy such “super-Turing” capabilities, particularly that of Penrose. One more is given below.

The second sees no reason for nature to submit to the constrictions imposed on it by computability theory. Kripke [81, p. 91] writes:

Consider such dimensionless physical constants as the electron-proton mass ratio, or the fine structure constant. As far as I have heard, at the present time physical theory has nothing much to say about the mathematical properties of these numbers, whether they are algebraic or transcendental, or even rational or irrational. In particular, it might well be the case that the decimal expansions of these numbers are not Turing computable. Assume this is so, and also assume that time is infinite in extent and that there are no limitations in energy that prevent computations of these decimals to any number of places. As far as I can see, nothing in principle rules this out. Then an empirically possible machine would exist that calculates a decimal not computable in Turing’s sense.

Many arguments for and against hypercomputation are collected in [129].

6.2.1 Existing Devices

Biological Devices

Allowing one to compute with the aid of any physical device and taking biology as an aspect of physics, we end up with having humans on both sides of the computation – both as those who perform the computation using a physical-biological device and as the device itself. With such an outlook, the border between the physical Church-Turing thesis and the AI Thesis gets even blurrier.

Bringsjord’s Arguments

Selmer Bringsjord and his colleagues claim to have mustered strong evidence that human beings can compute more than Turing machines. They claim that people can solve the “busy beaver” problem, which is known to be beyond the capabilities of Turing machines. This is intended as an attack on computationalism, along the lines of Gödel’s reflections.⁶

Bringsjord embarked on an NSF-funded project to compute explicit values of the Busy Beaver function, called here Σ^B . Given its success in evaluating several values, he then proceeds to make the following claim:

Persons must in some real way be capable of infinitary information processing. . . .

The idea is really quite simple: If humans are smart enough to determine $\Sigma^B(n)$, they will eventually (perhaps after 100 years, perhaps after 1000, perhaps after 1,000,000,000, . . .) be smart enough to determine $\Sigma^B(n + 1)$ While [this premise] may not be unassailable,

⁶ Another argument of Bringsjord, explicitly touted as “a case against Church’s thesis”, takes its cue from the perceived ability of humans to judge the literary quality of writings that they cannot for the life of them produce themselves [20].

a variant would seem to be, viz., that if humans are smart enough to determine $\Sigma^B(n)$, it's then mathematically possible that they determine $\Sigma^B(n+1)$. [21]

Owen Kellett, Bringsjord's student and one of the authors of [21], elaborates:

Human visual reasoning system possesses [natural, physical] computational powers beyond the Turing limit, and these powers can be employed in efforts to solve $\Sigma(n)$ The Busy Beaver function . . . is still computable by the fantastic capabilities of the human mind.

While Turing machines hold their place in conventional computability theory, the laws of the physical world are beyond the computational expressive power of Turing machines. Therefore, humans, a part of this physical universe, can harness this power to perform computations beyond the Turing limit. [71]

Thus, this argument is akin to Penrose's, appealing to unknown natural hypercomputational mechanisms that are somehow harnessed by biological brains. As such it is an attack on the physical thesis (or the non-mathematical one). In his book with Michael Zenzen, *Superminds*, Bringsjord summarizes his point of view:

Persons are superminds: they are simple souls capable of cognition capturable by standard computationalism, cognition capturable by *hypercomputationalism*, and cognition that will forever resist capturing in any third-person scheme. [22, §7.3]

Bringsjord conveniently lists the main objections to this claim, presumably culled from reviews and other reactions:

- This assertion resembles the claim that the human record for the 100 meter run will converge to 0 seconds, since we achieve a new record with each passing year. . . . Solving the busy beaver problem for input up to 6 . . . doesn't tell us much about the limits of the human mind.⁷
- The theorem . . . only implies that humans will get stuck billions of orders of magnitude past where current research on computing the function stands.
- Humans are far from immune to errors.
- There is some confusion of a specific person with all of humanity, past, present, and future.
- There is . . . a hidden assumption that people are unaffected by nature. . . . Historically, many scientific breakthroughs happen serendipitously.

Randomness

True randomness is believed by most physicists to exist in nature [127], whereas an ordinary Turing machine is incapable of generating truly random values. Hence, we simply get that random sequences, which can be computed by actual devices, falsify the physical thesis [19, p. 74], [24, p. 10].

⁷ Herbert Robbins, the statistician: "Nobody is going to run 100 meters in five seconds, no matter how much is invested in training and machines. The same can be said about using the brain. The human mind is no different now from what it was five thousand years ago. And when it comes to mathematics, you must realize that this is the human mind at an extreme limit of its capacity." [92]

A counterargument is that a *random number* (qua indexed sequence of digits) is not a calculation in the intended sense of the thesis. Proper functions always provide the identical output to the same input, which is not the case if a device randomly generates output values or uses such values to guide computation.

The above remarks converge on the conclusion that genuine random processes are not computations. Unlike computations properly so-called, random processes cannot be used to generate the desired values of a function or solve the desired instances of a general problem. [97, p. 751]

Quantum physicist, David Deutsch, writes:

Quantum computing machines, and indeed classical stochastic computing machines, do not “compute functions” in the above sense: the output state of a stochastic machine is random with only the probability distribution function for the possible outputs depending on the input state. [44]

So we might ask instead that the distribution of outputs satisfy some requirements.

In cases where a random process provides the same one output with high probability, it can be considered an *estimation* of a proper function, and we may ask whether the function that it estimates is computable by a Turing machine. The answer to this question is “yes”, as one may derandomize such a process, trying all paths until some required percentage give the same answer [55] (though this does depend on the precise definitions [41]).

A different way to handle random numbers in a “proper function” setting might be to fix a computation by augmenting it with a memory cache. Consider a device M that maintains an unbounded array R of binary values, initially blank. Whenever M is presented with an input $n \in \mathbb{N}$, it does the following: If $R[n]$ is blank, it randomly chooses between 0 and 1, stores this bit in $R[n]$, and returns it as the value of the function on n ; on the other hand, if $R[n]$ already has value v , it returns that v . Such a device computes a proper function $R : \mathbb{N} \rightarrow \{0, 1\}$, but has the characteristics of an *interactive* process, which in itself need not make it ineffective or non-Turing-computable, as discussed in Section 2. Yet, because it *learns* and *evolves* with time, one may argue whether it is better viewed as an infinite series of distinct computations, rather than as a device for repeatedly computing a single function. An additional reason to contend that this is not a valid attack on the thesis is the uniqueness of the device – any attempt to clone it will result in a different device, with different behavior, computing a different function.

Evolution and Learnability

An interactive process evolves over time – it maintains a memory that may grow and change from one invocation of the process to another. A central question is then how does it evolve: Is its evolution predefined or not?

There are claims against the physical thesis, arguing that internet applications, such as Google Translate, transcend Turing computability (e.g. [64]). So as to consider such an evolving process as a proper function that always returns the same

output for the same input, one may augment every “genuine” call to the application with a unique number and cache all previous queries as explained above with respect to a random process.

The main argument against such claims is that the evolution of the application is not predefined. Google Translate, for example, changes over time not as a result of a predefined learning algorithm, but as a result of humans modifying it. Hence, claims that such computations exceed Turing computability are analogous to claims that the human mind, and possibly all of future humanity, exceed Turing computability in the limit. Such processes are usually not considered to lie under the rubric of effective means [76, p. 50].

6.2.2 Futuristic Devices

We list below some of the claims about theoretical devices that may agree with modern physical theories and compute more than Turing machines.

Zeno Machines

A *Zeno machine* or *accelerating machine* can be thought of as a standard Turing machine in which each step is performed in, say, half the time of the previous step [138, 57, 102, 34]. Having such an accelerating machine, one can solve the halting problem: If the first step takes, for example, a second, any number of steps takes less than two seconds. Hence, after two seconds we (the machine users) can inspect the machine’s tape and check if it has a halting mark. In case that it does not, we know that it does not halt. The logical possibility of such a device has long been debated [61, 34]. On the other hand, there are claims that such an accelerating device, as well as communication with it, is theoretically possible in the vicinity of certain kinds of black holes by certain interpretations of special and general relativity [98, 68, 67, 48, 113].

One criticism of this view is that its physical assumptions do not match current physical theories and that, even under these assumptions, communication is only possible from Earth to the vicinity of the black hole, and not the other way round, providing a “one-way ticket” to knowing whether computation has ended [45, 52], as well as logical considerations [36].

Accessible Reals

Equipped somehow with an infinite precision real number, a device can easily compute beyond Turing machines. A simple example is Fred Abramson’s Extended Turing Machine, which can store a real number on a single square of its tape [2].

A more classic example of calculating with reals is ruler (straightedge) and compass (*RC*) computations. A common algorithmic definition runs something like this:

Given a finite set of points $\mathcal{B} = \{B_0, \dots, B_m\}$ in the Euclidean plane, a point P is *RC-constructible from the set \mathcal{B}* if there is a finite set of points $\{P_0, \dots, P_n\}$ such that $P = P_n$, $P_0 \in \mathcal{B}$ and every point P_i ($1 \leq P_i \leq n$) is either a point of \mathcal{B} or is at the intersection either of two lines, or of a line and a circle, or of two circles, themselves obtained as follows:

- any considered line passes through two distinct points from the set $\{P_0, \dots, P_{i-1}\}$;
- any considered circle has its center in the set $\{P_0, \dots, P_{i-1}\}$ and its radius is equal to the distance $P_j P_k$ for some $j < i$ and $k < i$. [26, trans. [109]]

The widespread view is that *RC*-computations do not exceed Turing computability, merely allowing one to compute a certain family of algebraic numbers (Wantzel’s Theorem; see [85]). Yet, if initial points are arbitrary real points in the plane – given either as inputs or as elements of the initial state of a device, then with the addition of a “unit” segment that represents the natural number 1 and basic geometric operations, one can compute non-Turing-computable functions over the natural numbers [6].

A more involved example is Hava Siegelmann’s version of neural networks [124].

The main criticism of these types of computation, from an effectiveness point of view, is the inaccessibility of the infinite precision needed to set up the computation framework [32], and that “the claimed noncomputability was nothing more than that of the real numbers provided at the beginning” [39]. But were there a primordial uncomputable fundamental constant in the universe – and *a priori* there is no physical reason to preclude that possibility, then accessing its digits by means of some idealized experiment might serve to violate the physical thesis, as already noted [81, p. 91].

Quantum Computers

The dominant view of quantum computation is that, if it is feasible and scalable, it will allow one to compute faster than Turing machines but will not enable one to compute more [44, 3]. If this view holds true, it would disprove the “extended” thesis, but not the original one. Still, there are also more audacious claims that quantum mechanics might allow for computing devices that exceed Turing computability [25, 91, 72]. Tien Kieu, for example, suggests a way to use the Quantum Adiabatic Theorem to provide a computational solution to Hilbert’s Tenth Problem, which is known to be Turing uncomputable [72]. Criticisms of Kieu’s proposal include issues of representation and concerns that adiabatic evolution fails to actually reach the required solution [39, 132, 125, 65].

Deutsch’s Arguments

Another reaction to the stance that natural random processes falsify the physical thesis is to assert that computability at its core is indeed Turing-machine computability.

Having said that, to extend the Turing model's scope from classical functions to stochastic functions, one merely needs to augment it with a (physical or oracular) source of randomness. This, as we have seen, accords with Turing's own perspective when describing the "choice machine" variant of his eponymous machines: "When such a machine reaches one of these ambiguous configurations, it cannot go on until some arbitrary choice has been made by external operator" [133]. In this context, the "external operator" would be a physical apparatus, rather than a human guide.

Deutsch reformulated the physical Church-Turing thesis as an explicit postulate about nature herself, what is now known as "Deutsch's Principle":

Every finitely realizable physical system can be perfectly simulated by a universal model computing machine operating by finite means. [44]

Specifically, he posits that a *quantum* generalization of the Turing machine can simulate every *physical* process, in the sense that for every possible input state, represented in some fixed way, the distribution of outcomes of output experiments is the same (up to representation) as for the physical system:

A class of model computing machines that is the quantum generalization of the class of Turing machines is described, and it is shown that quantum theory and the 'universal quantum computer' are compatible with the principle. [44]

Empowered this way, Deutsch's proposed quantum machine computes discrete functions faster than Turing's machine can, can compute stochastic outputs that no classical system can reproduce, and can compute continuous functions to arbitrarily high (albeit imperfect) accuracy [44].

It is, of course, an open (unanswerable) question whether there is at all any complete, immutable, finite description of the (observable) behavior of the physical universe [27].

7 Conclusion

In summary, most computer scientists take the mathematical thesis for granted, based on the strong arguments put forth by Turing and those who followed him.

Moreover, there is a long list of highly regarded logicians who are of the opinion that there is room for formalization and perhaps proof of the mathematical thesis, including: Gödel, Gandy, Mendelson, Shoenfield, Shapiro, George Kreisel [79], Sieg, Kripke, Soare, Robert Black [11], Friedman, and Gurevich. Some have even proposed formal proofs or sketches of proofs. Other scholars, starting with Church, Turing, and Kleene and continuing to the present, maintain that there is an impassable dissonance between intuitive and formal notions of effectiveness.

The physical thesis, on the other hand, enjoys far less support, as also purported claims of hypercomputability.

Acknowledgments

We are enormously grateful for comments and suggestions from Arnon Avron, Erwin Engeler, Oron Shagrir, Wilfried Sieg, and an anonymous reader.

References

1. Aaronson, S.: The toaster-enhanced Turing machine (2012). URL <https://www.scottaaronson.com/blog/?p=1121>. The Blog of Scott Aaronson
2. Abramson, F.G.: Effective computation over the real numbers. In: 12th Annual Symposium on Switching and Automata Theory (SWAT 1971), pp. 33–37 (1971)
3. Arrighi, P., Dowek, G.: The physical Church-Turing thesis and the principles of quantum theory. *International Journal of Foundations of Computer Science* **23**(5), 1131–1145 (2012)
4. Arrighi, P., Dowek, G.: The principle of a finite density of information. In: H. Zenil (ed.) *Irreducibility and Computational Equivalence, Emergence, Complexity and Computation*, vol. 2, pp. 127–134. Springer, Berlin (2013). DOI 10.1007/978-3-642-35482-3_11
5. Avron, A.: Mishpete Gedel u-ve‘ayat ha-yesodot shel ha-matematikah (= Gödel’s Theorems and the Problem of the Foundations of Mathematics). Broadcast University, Ministry of Defence, Jerusalem, Israel (1998). In Hebrew
6. Avron, A.: Personal communication (2020)
7. Avron, A.: The problematic nature of Gödel’s disjunctions and Lucas-Penrose’s theses. *Semiotic Studies* (to appear)
8. Barendregt, H.: The impact of the lambda calculus in logic and computer science. *Bulletin of Symbolic Logic* **3**(2), 181–215 (1997). URL <http://www.math.ucla.edu/~asl/bsl/0302/0302-003.ps>
9. Barwise, J.: An introduction to first-order logic. In: J. Barwise (ed.) *Handbook of Mathematical Logic*, chap. A.1, pp. 5–46. North-Holland (1977)
10. Beckman, D., Gottesman, D., Nielsen, M.A., Preskill, J.: Causal and localizable quantum operations. *Phys. Rev. A* **64** (2001)
11. Black, R.: Proving Church’s Thesis. *Philosophia Mathematica* **8**(3), 244–258 (2000). DOI 10.1093/phimat/8.3.244
12. Boker, U., Dershowitz, N.: Abstract effective models. In: M. Fernández, I. Mackie (eds.) *New Developments in Computational Models: Proceedings of the First International Workshop on Developments in Computational Models (DCM 2005)*, Lisbon, Portugal (July 2005), *Electronic Notes in Theoretical Computer Science*, vol. 135, pp. 15–23 (2006)
13. Boker, U., Dershowitz, N.: Comparing computational power. *Logic Journal of the IGPL* **14**(5), 633–648 (2006). URL <http://nachum.org/papers/ComparingComputationalPower.pdf>
14. Boker, U., Dershowitz, N.: A hypercomputational alien. *Applied Mathematics and Computation* **178**(1), 44–57 (2006)
15. Boker, U., Dershowitz, N.: The Church-Turing thesis over arbitrary domains. In: A. Avron, N. Dershowitz, A. Rabinovich (eds.) *Pillars of Computer Science, Essays Dedicated to Boris (Boaz) Trakhtenbrot on the Occasion of His 85th Birthday*, *Lecture Notes in Computer Science*, vol. 4800, pp. 199–229. Springer (2008). URL <http://nachum.org/papers/ArbitraryDomains.pdf>
16. Boker, U., Dershowitz, N.: The influence of domain interpretations on computational models. *Journal of Applied Mathematics and Computation* **215**(4), 1323–1339 (2009)
17. Boker, U., Dershowitz, N.: Three paths to effectiveness. In: A. Blass, N. Dershowitz, W. Reisig (eds.) *Fields of Logic and Computation: Essays Dedicated to Yuri Gurevich on the Occasion of His 70th Birthday*, *Lecture Notes in Computer Science*, vol. 6300, pp. 36–47. Springer, Berlin (2010). URL <http://nachum.org/papers/ThreePathsToEffectiveness.pdf>

18. Bournez, O., Dershowitz, N., Néron, P.: An axiomatization of analog algorithms. In: Computability in Europe 2016: Pursuit of the Universal (CiE, Paris, France), *Lecture Notes in Computer Science*, vol. 9709, pp. 215–224. Springer, Switzerland (2016). URL <http://nachum.org/papers/AxiomatizationAnalog.pdf>. Full version at <https://arxiv.org/pdf/1604.04295v2.pdf>
19. Bowie, G.L.: An argument against Church's thesis. *The Journal of Philosophy* **70**(3), 66–76 (1973)
20. Bringsjord, S., Ferrucci, D.A.: The narrative-based refutation of Church's thesis. In: *Artificial Intelligence and Literary Creativity: Inside the Mind of BRUTUS, a Storytelling Machine*, chap. 5, pp. 105–148. Lawrence Erlbaum, Mahwah, NJ (2000)
21. Bringsjord, S., Kellett, O., Shilliday, A., Taylor, J., van Heuveln, B., Yang, Y., Baumes, J., Ross, K.: A new Gödelian argument for hypercomputing minds based on the Busy Beaver problem. *Applied Mathematics and Computation* **176**(2), 516–530 (2006). DOI <https://doi.org/10.1016/j.amc.2005.09.071>
22. Bringsjord, S., Zenzen, M.: Superminds: People Harness Hypercomputation, and More, *Studies in Cognitive Systems*, vol. 29. Kluwer Academic, Dordrecht, The Netherlands (2003)
23. Broy, M.: Computability and realizability for interactive computations. *Information and Computation* **241**, 277–301 (2015)
24. Calude, C.S.: Algorithmic randomness, quantum physics, and incompleteness. In: *Proceedings of the Conference on Machines, Computations and Universality (MCU 2004)*, vol. 3354, pp. 1–17 (2004)
25. Calude, C.S., Pavlov, B.: Coins, quantum measurements, and Turing's barrier. *Quantum Information Processing* **1**(1), 107–127 (2002)
26. Carréga, J.C.: *Théorie des corps - La règle et le compas*. Hermann, Paris (1981)
27. Carroll, J.W.: Laws of nature. In: E.N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy*, fall 2016 edn. Metaphysics Research Lab, Stanford University (2016). URL <https://plato.stanford.edu/archives/fall2016/entries/laws-of-nature>
28. Chalmers, D. (ed.): *Symposium on Roger Penrose's Shadows of the Mind*, vol. 2. Association for the Scientific Study of Consciousness (1995). URL <http://journalpsyche.org/archive/volume-2-1995-1996>
29. Chalmers, D.: A computational foundation for the study of cognition. *Journal of Cognitive Science* **12**(4), 325–359 (2011). DOI 10.17791/jcs.2011.12.4.325. Written in 1993
30. Church, A.: An unsolvable problem of elementary number theory. *American Journal of Mathematics* **58**, 345–363 (1936)
31. Church, A.: Review of “On computable numbers, with an application to the Entscheidungsproblem”. *The Journal of Symbolic Logic* **2**(1), 42–43 (1937)
32. Cockshotta, P., Mackenzie, L., Michaelson, G.: Physical constraints on hypercomputation. *Theoretical Computer Science* **394**, 159–174 (2008)
33. Copeland, B., Proudfoot, D.: Deviant encodings and Turing's analysis of computability. *Studies in History and Philosophy of Science* **41**, 247–252 (2010). DOI 10.1016/j.shpsa.2010.07.010
34. Copeland, B.J.: Accelerating Turing machines. *Minds and Machines* **12**(2), 281–300 (2002). DOI 10.1023/A:1015607401307
35. Copeland, J.B., Shagrir, O.: Physical computation: How general are Gandy's principles for mechanisms? *Minds and Machines* **17**(2), 217–231 (2007)
36. Cotogno, P.: Hypercomputation and the physical Church-Turing thesis. *The British Journal for the Philosophy of Science* **54**(2), 181–223 (2003)
37. Davis, M.: Why Gödel didn't have Church's Thesis. *Information and Control* **54**(1/2), 3–24 (1982)
38. Davis, M.: The myth of hypercomputation. In: C. Teuscher (ed.) *Alan Turing: Life and Legacy of a Great Thinker*, pp. 195–212. Springer (2003)
39. Davis, M.: The Church-Turing thesis: Consensus and opposition. In: *Logical Approaches to Computational Barriers*, pp. 125–132. Springer, Berlin (2006)
40. Davis, M.: *The Universal Computer: The Road from Leibniz to Turing*, 3rd edn. CRC Press (2018)

41. de Leeuw, K., Moore, E.F., Shannon, C.E., Shapiro, N.: Computability by probabilistic machines. In: C.E. Shannon, J. McCarthy (eds.) *Automata Studies, Annals of Mathematics Studies*, vol. 34, pp. 183–212. Princeton University Press (1956)
42. Dershowitz, N.: The four sons of Penrose. In: Proceedings 12th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR) (Montego Bay, Jamaica), *Lecture Notes in Computer Science*, vol. 3835, pp. 125–138. Springer (2005)
43. Dershowitz, N., Gurevich, Y.: A natural axiomatization of computability and proof of Church's Thesis. *Bulletin of Symbolic Logic* **14**(3), 299–350 (2008)
44. Deutsch, D.: Quantum theory, the Church-Turing principle and the universal quantum computer. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **400**, 97–117 (1985)
45. Earman, J., Norton, J.D.: Forever is a day: Supertasks in Pitowsky and Malament-Hogarth spacetimes. *Philosophy of Science* **60**(1), 22–42 (1993)
46. van Emde Boas, P.: Machine models and simulations. In: J. van Leeuwen (ed.) *Handbook of Theoretical Computer Science*, vol. A: Algorithms and Complexity, pp. 1–66. North-Holland, Amsterdam (1990). URL <http://www.illc.uva.nl/Research/Reports/CT-1988-05.text.pdf>
47. Engeler, E.: *The Combinatory Programme*. Progress in Theoretical Computer Science. Birkhäuser, Boston (1995)
48. Etesi, G., Németi, I.: Non-Turing computations via Malament-Hogarth space-times. *International Journal of Theoretical Physics* **41**(2), 341–370 (2002)
49. Feferman, S.: Theses for computation and recursion on concrete and abstract structures. In: *Turing's Revolution: The Impact of His Ideas about Computability*, pp. 105–126. Springer International Publishing (2015)
50. Folina, J.: Church's thesis: Prelude to a proof. *Philosophia Mathematica* **6**(3), 302–323 (1998)
51. Friedman, H.M.: Mathematical logic in the 20th and 21st centuries (2000). URL <http://cs.nyu.edu/pipermail/fom/2000-April/003913.html>. FOM mailing list. April 27, 2000
52. Galton, A.: The Church-Turing thesis: Still valid after all these years? *Applied Mathematics and Computation* **178**, 93–102 (2006)
53. Gandy, R.: Church's thesis and principles for mechanisms. In: J. Barwise, D. Kaplan, H.J. Keisler, P. Suppes, A.S. Troelstra (eds.) *The Kleene Symposium, Studies in Logic and The Foundations of Mathematics*, vol. 101, pp. 123–148. North-Holland (1980)
54. Gandy, R.: The confluence of ideas in 1936. In: *A Half-Century Survey on the Universal Turing Machine*, pp. 55–111. Oxford University Press, Inc., New York, NY, USA (1988). URL <http://dl.acm.org/citation.cfm?id=57249.57252>
55. Gill, J.: Computational complexity of probabilistic turing machines. *SIAM Journal on Computing* **6**(4), 675–695 (1977)
56. Gödel, K.: Some remarks on the undecidability results. In: S. Feferman, J. Dawson, S. Kleene (eds.) *Kurt Gödel: Collected Works, Vol. II*, pp. 305–306. Oxford University Press (1972)
57. Gold, E.M.: Limiting recursion. *J. Symbolic Logic* **30**(1), 28–48 (1965). DOI 10.2307/2270580
58. Goldin, D., Wegner, P.: The Church-Turing Thesis: Breaking the myth. In: S.B. Cooper, B. Löwe, L. Torenlvliet (eds.) *New Computational Paradigms: First Conference on Computability in Europe (CiE 2005, Amsterdam), Lecture Notes in Computer Science*, vol. 3526, pp. 152–168. Berlin (2005)
59. Goldin, D., Wegner, P.: The interactive nature of computing: Refuting the strong church-turing thesis. *Minds and Machines* **18**, 17–38 (2008)
60. Goldin, D.Q., Smolka, S.A., Wegner, P. (eds.): *Interactive Computation: The New Paradigm*. Springer, Berlin (2006)
61. Grünbaum, A.: Messrs. Black and Taylor on temporal paradoxes. *Analysis* **12**(6), 144–148 (1952). URL <http://www.jstor.org/stable/3326977>
62. Gurevich, Y.: Sequential abstract state machines capture sequential algorithms. *ACM Transactions on Computational Logic* **1**, 77–111 (2000)

63. Gurevich, Y.: What is an algorithm? In: SOFSEM 2012: Theory and Practice of Computer Science, pp. 31–42 (2012)
64. Gurevich, Y.: Unconstrained Church-Turing thesis cannot possibly be true. *Bull. EATCS* **127** (2019). URL <http://bulletin.eatcs.org/index.php/beatcs/article/view/566/565>
65. Hodges, A.: Can quantum computing solve classically unsolvable problems? arXiv (2005). URL <http://arXiv.org/abs/quant-ph/0512248>
66. Hodges, A.: Did Church and Turing have a thesis about machines? In: A. Olszewski, J. Wolenski, R. Janusz (eds.) *Church's Thesis After 70 Years*, pp. 214–224. Ontos Verlag (2006)
67. Hogarth, M.: Non-Turing computers and non-Turing computability. In: *Proceedings of the Philosophy of Science Association*, vol. 1994, pp. 126–138 (1994)
68. Hogarth, M.L.: Does general relativity allow an observer to view an eternity in a finite time? *Foundations of Physics Letters* **5**(2), 173–181 (1992)
69. Hopcroft, J.E., Ullman, J.D.: *Formal Languages and Their Relation to Automata*. Addison-Wesley, Reading, MA (1968)
70. Kalmár, L.: An argument against the plausibility of Church's thesis. In: A. Heyting (ed.) *Constructivity in Mathematics, Proceedings of the Colloquium Held at Amsterdam, 1957*, pp. 72–80. North-Holland, Amsterdam (1959)
71. Kellett, O.: A multi-faceted attack on the Busy Beaver problem. Master's thesis, Rensselaer Polytechnic Institute, Troy, New York (2005)
72. Kieu, T.D.: Quantum algorithm for Hilbert's Tenth Problem. *International Journal of Theoretical Physics* **42**, 1461–1478 (2003)
73. Kleene, S.C.: Lambda-definability and recursiveness. *Duke Mathematical Journal* **2**, 340–353 (1936)
74. Kleene, S.C.: *Introduction to Metamathematics*. North Holland (1952)
75. Kleene, S.C.: Origins of recursive function theory. *Annals of the History of Computing* **3**(1), 52–67 (1981)
76. Kleene, S.C.: Turing's analysis of computability, and major applications of it. In: *A Half-century Survey on The Universal Turing Machine*, pp. 17–54. Oxford University Press, Inc. (1988)
77. Knuth, D.E.: Algorithm and program; information and data. *Communications of the ACM* **9**, 654 (1968)
78. Kolmogorov, A.N.: O ponyatii algoritma [On the concept of algorithm] (in Russian). *Uspekhi Matematicheskikh Nauk [Russian Mathematical Surveys]* **8**(4), 1175–1176 (1953). English version in: Vladimir A. Uspensky and Alexei L. Semenov, *Algorithms: Main Ideas and Applications*, Kluwer, Norwell, MA, 1993, pp. 18–19
79. Kreisel, G.: Mathematical logic: What has it done for the philosophy of mathematics? In: R. Schoenman (ed.) *Bertrand Russell: Philosopher of the Century*. George Allen and Unwin, London (1967)
80. Kripke, S.A.: The origins and nature of computation (2006). URL <https://www.youtube.com/watch?v=D9SP5wj882w>. Presented at the 21st International Workshop on the History and Philosophy of Science
81. Kripke, S.A.: The Church-Turing “thesis” as a special corollary of Gödel's completeness theorem. In: J. Copeland, C. Posy, O. Shagrir (eds.) *Computability: Turing, Gödel, Church, and Beyond*, pp. 77–104. MIT Press (2013)
82. LaForte, G., Hayes, P.J., Ford, K.M.: Why Gödel's theorem cannot refute computationalism. *Artificial Intelligence* **104**(1–2), 265–286 (1998)
83. Lewis, H.R., Papadimitriou, C.H.: *Elements of the Theory of Computation*. Prentice-Hall, Englewood Cliffs, NJ (1981)
84. Lucas, J.R.: Minds, machines and Gödel. *Philosophy* **XXXVI**, 112–127 (1961). URL <http://users.ox.ac.uk/~jr1lucas/mmg.html>. Reprinted in *The Modeling of Mind*, K. M. Sayre and F. J. Crosson, eds., Notre Dame Press, 1963, pp. 269–270
85. Lützen, J.: Why was Wantzel overlooked for a century? The changing importance of an impossibility result. *Historia Mathematica* **36**(4), 374–394 (2009). DOI 10.1016/j.hm.2009.03.001

86. Mendelson, E.: On some recent criticism of Church's thesis. *Notre Dame Journal of Formal Logic* **IV**(3), 201–205 (1963)
87. Mendelson, E.: Second thoughts about Church's thesis and mathematical proofs. *Journal of Philosophy* **87**(5), 225–233 (1990)
88. Mendelson, E.: *Introduction to Mathematical Logic*, 5th edn. *Discrete Mathematics and Its Applications*. CRC Press (2009). URL <https://books.google.co.il/books?id=NJ7SBQAAQBAJ>
89. Moschovakis, Y.N.: Review of four recent papers on Church's thesis. *Journal of Symbolic Logic* **33**(3), 471–472 (1968). DOI 10.2307/2270346
90. Moschovakis, Y.N.: What is an algorithm? In: B. Engquist, W. Schmid (eds.) *Mathematics Unlimited — 2001 and Beyond*, pp. 919–936. Springer, Berlin (2001)
91. Nielsen, M.A.: Computable functions, quantum measurements, and quantum dynamics. *Physical Review Letters* **79**(15), 2915–2918 (1997)
92. Page, W.: An interview with Herbert Robbins. *The Two-Year College Mathematics Journal* **15**(1), 2–24 (1984). DOI 10.1080/00494925.1984.11972739
93. Parberry, I.: Parallel speedup of sequential machines: A defense of parallel computation thesis. *SIGACT News* **18**(1), 54–67 (1986). DOI 10.1145/8312.8317
94. Penrose, R.: *The Emperor's New Mind: Concerning Computers, Minds, and The Laws of Physics*. Oxford University Press, New York (1989)
95. Penrose, R.: *Shadows of the Mind: A Search for the Missing Science of Consciousness*. Oxford University Press, Oxford (1994)
96. Péter, R.: Rekursivität und konstruktivität. In: A. Heyting (ed.) *Constructivity in Mathematics*, *Proceedings of the Colloquium Held at Amsterdam, 1957*, pp. 226–233. North-Holland, Amsterdam (1959)
97. Piccinini, G.: The physical Church-Turing thesis: Modest or bold? *The British Journal for the Philosophy of Science* **62**, 733–769 (2011)
98. Pitowsky, I.: The physical Church thesis and physical computational complexity. *Iyyun: The Jerusalem Philosophical Quarterly* **39**, 81–99 (1990)
99. Porte, J.: Quelques pseudo-paradoxes de la 'calculabilité effective'. In: *Actes du 2me Congrès International de Cybernétique*, pp. 332–334. Namur, Belgium (1960)
100. Post, E.L.: Finite combinatory processes—formulation I. *Journal of Symbolic Logic* **1**(3), 103–105 (1936)
101. Post, E.L.: Absolutely unsolvable problems and relatively undecidable propositions: Account of an anticipation. In: M. Davis (ed.) *Solvability, Provability, Definability: The Collected Works of Emil L. Post*, pp. 375–441. Birkhäuser, Boston, MA (1994). Unpublished notes, 1941
102. Putnam, H.: Trial and error predicates and the solution to a problem of Mostowski. *J. Symbolic Logic* **30**(1), 49–57 (1965). DOI 10.2307/2270581
103. Putnam, H.: Book review: *Shadows of the Mind* by Roger Penrose. *Bulletin of the American Mathematical Society* **32**(3), 370–373 (1995)
104. Rapaport, W.J.: *Philosophy of Computer Science*. Online draft (2020). URL <https://cse.buffalo.edu/~rapaport/Papers/physics.pdf>
105. Reisig, W.: The computable kernel of Abstract State Machines. *Theoretical Computer Science* **409**, 126–136 (2008)
106. Rescorla, M.: Church's thesis and the conceptual analysis of computability. *Notre Dame Journal of Formal Logic* **48**(2), 253–280 (2007)
107. Rescorla, M.: Copeland and Proudfoot on computability. *Studies in History and Philosophy of Science Part A* **43**(1), 199–202 (2012). DOI 10.1016/j.shpsa.2011.11.002. Reconsidering the Dynamics of Reason: A Symposium in Honour of Michael Friedman
108. Rogers Jr., H.: *Theory of Recursive Functions and Effective Computability*. McGraw-Hill, New York (1966)
109. Schreck, P.: On the mechanization of straightedge and compass constructions. *Journal of Systems Science and Complexity* **32**, 124–149 (2019)
110. Schumacher, B., Westmoreland, M.D.: Locality and information transfer in quantum operations. *Quantum Information Processing* **4**(1), 13–34 (2005)

111. Shagrir, O.: Effective computation by humans and machines. *Minds and Machines* **12**, 221–240 (2002)
112. Shagrir, O.: Gödel on Turing on computability. In: A. Olszewski, J. Wolenski, R. Janusz (eds.) *Church's Thesis after 70 Years*, pp. 393–419. Ontos-Verlag (2006)
113. Shagrir, O., Pitowsky, I.: Physical hypercomputation and the Church-Turing thesis. *Minds and Machines* **13**(1), 87–101 (2003)
114. Shapiro, S.: Acceptable notation. *Notre Dame Journal of Formal Logic* **23**(1), 14–20 (1982)
115. Shapiro, S.: Understanding Church's thesis, again. *Acta Analytica* **11**, 59–77 (1993)
116. Shapiro, S.: Proving things about the informal. In: G. Sommaruga, T. Strahm (eds.) *Turing's Revolution: The Impact of his Ideas About Computability*, pp. 283–296. Springer, Cham (2015). DOI 10.1007/978-3-319-22156-4_11
117. Shepherdson, J.C.: On the definition of computable function of a real variable. *Zeitschrift für mathematische Logik und Grundlagen der Mathematik (Mathematical Logic Quarterly)* **22**(1), 391–402 (1976). DOI 10.1002/malq.19760220148
118. Shoenfield, J.R.: *Recursion Theory, Lecture Notes In Logic*, vol. 1. Springer-Verlag, Heidelberg, New York (1991)
119. Sieg, W.: Mechanical procedures and mathematical experience. In: A. George (ed.) *Mathematics and Mind*, pp. 71–117. Oxford University Press (1994)
120. Sieg, W.: Step by recursive step: Church's analysis of effective calculability. *Bulletin of Symbolic Logic* **3**, 154–180 (1997). DOI 10.2307/421012
121. Sieg, W.: Church without dogma: Axioms for computability. In: S.B. Cooper, B. Löwe, A. Sorbi (eds.) *New Computational Paradigms: Changing Conceptions of What is Computable*, pp. 139–152. Springer, New York (2007)
122. Sieg, W.: Axioms for computability: Do they allow a proof of Church's thesis? In: H. Zenil (ed.) *A Computable Universe. Understanding and Exploring Nature as Computation*, pp. 99–123. World Scientific/Imperial College Press, Singapore (2013)
123. Sieg, W., Byrnes, J.: An abstract model for parallel computations: Gandy's thesis. *The Monist* **82**(1), 150–164 (1999)
124. Siegelmann, H.T.: *Neural Networks and Analog Computation: Beyond the Turing Limit*. Birkhäuser, Boston (1998)
125. Smith, W.D.: Three counterexamples refuting Kieu's plan for "quantum adiabatic hypercomputation"; and some uncomputable quantum mechanical tasks. *Applied Mathematics and Computation* **178**(1), 184–193 (2006). Special Issue on Hypercomputation
126. Soare, R.I.: The history and concept of computability. In: *Handbook of Computability Theory*, pp. 3–36. Elsevier (1999)
127. Stipčević, M., Koç, Ç.K.: True random number generators. In: Ç.K. Koç (ed.) *Open Problems in Mathematics and Computational Science*, pp. 275–315. Springer International Publishing (2014)
128. Strachey, C.: An impossible program. *Computer Journal* **7**(4), 313 (1965). DOI 10.1093/comjnl/7.4.313
129. Syropoulos, A.: *Hypercomputation: Computing Beyond the Church-Turing Barrier*. Springer Science & Business Media (2008)
130. Szabó, M.: Kalmár's argument against the plausibility of Church's thesis. *History and Philosophy of Logic* **39**(2), 140–157 (2018). DOI 10.1080/01445340.2017.1396520
131. Takahashi, M.: A foundation of finite mathematics. *Publications of the Research Institute for Mathematical Sciences* **12**, 577–708 (1977)
132. Tsirelson, B.: The quantum algorithm of Kieu does not solve the Hilbert's tenth problem. arXiv (2001). URL <http://arXiv.org/abs/quant-ph/0111009>
133. Turing, A.M.: On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society* **42**, 230–265 (1936–37). URL <http://www.abelard.org/turpap2/tp2-ie.asp>. Corrections in vol. 43 (1937), pp. 544–546. Reprinted in M. Davis (ed.), *The Undecidable*, Raven Press, Hewlett, NY, 1965
134. Wang, H.: *From Mathematics to Philosophy*. Keagan Paul, London, UK (1974)
135. Wang, H.: *Reflections on Kurt Gödel*. Bradford Books. MIT Press (1990)

136. Wang, H.: *A Logical Journey. From Gödel to Philosophy*. MIT Press (1996)
137. Wells, B.: Is there a nonrecursive decidable equational theory? *Minds and Machines* **12**(2), 301–324 (2002)
138. Weyl, H.: *Philosophy of Mathematics and Natural Science*. Princeton University Press (1949)
139. Wroclawski, M.: Representations of natural numbers and computability of various functions. In: F. Manea, B. Martin, D. Paulusma, G. Primiero (eds.) *Proceedings of the 15th Conference on Computability in Europe - Computing with Foresight and Industry (CiE 2019, Durham, UK)*, *Lecture Notes in Computer Science*, vol. 11558, pp. 298–309. Springer (2019). DOI 10.1007/978-3-030-22996-2_26