GUALTIERO PICCININI

COMPUTATIONALISM, THE CHURCH–TURING THESIS, AND THE CHURCH–TURING FALLACY

ABSTRACT. The Church–Turing Thesis (CTT) is often employed in arguments for computationalism. I scrutinize the most prominent of such arguments in light of recent work on CTT and argue that they are unsound. Although CTT does nothing to support computationalism, it is not irrelevant to it. By eliminating misunderstandings about the relationship between CTT and computationalism, we deepen our appreciation of computationalism as an empirical hypothesis.

Computationalism, or the Computational Theory of Mind, is the view that mental capacities are explained by inner computations. In the case of human beings, computationalists typically assume that inner computations are realized by neural processes; I will borrow a term from current neuroscience and refer to them as neural computations.[1] Typically, computationalists also maintain that neural computations are Turing-computable, that is, computable by Turing Machines (TMs). The Church–Turing thesis (CTT) says that a function is computable, in the intuitive sense, if and only if it is Turing-computable (Church 1936; Turing 1936–7). CTT entails that TMs, and any formalism equivalent to TMs, capture the intuitive notion of computation. In other words, according to CTT, if a function is computable in the intuitive sense, then there is a TM that computes it (or equivalently, it is Turing-computable).[2]

This applies to neural computations as well. Suppose that, as computationalism maintains, neural activity is computation, and suppose that the functions computed by neural mechanisms are computable in the intuitive sense. Then, by CTT, for any function computed by a neural mechanism, there is a TM that computes the same function. This is a legitimate argument for a technical version of computationalism, according to which neural computations are Turing-computable, from a generic one, according to which neural processes are computations in the intuitive sense, via CTT.

But should we believe CTT? The initial proponents of CTT, and most of CTT's supporters, appeal to a number of intuitive

Synthese (2007) 154: 97–120
DOI 10.1007/s11229-005-0194-z

considerations. The main considerations are that there are no known counterexamples, that various attempts at formalizing the notion of computation have yielded computationally equivalent formalisms, and that the notion of TM seems to capture well the intuitive notion of computation (Kleene 1952). Due to these considerations, most logicians and mathematicians believe CTT to be true. But some authors have attempted to go beyond intuitive considerations. They have attempted to establish (the technical version of) computationalism on independent grounds and use it to support CTT.

The first authors who employed the modern mathematical notion of computation to formulate a computational theory of mind and brain were McCulloch and Pitts (1943). Their theory entails that every neural activity is a computation (in the sense to which CTT is relevant), and that every mental activity is explained by some neural computation. The neural networks defined by McCulloch and Pitts are computationally less powerful than TMs, but if supplemented with a tape and the means to act on it, they are computationally equivalent to TMs. Since their neural networks are no more powerful than TMs, McCulloch and Pitts argue that their theory offers a "psychological justification" of CTT (ibid., p. 35). McCulloch and Pitts's argument for CTT rests on two premises: (i) that the intuitive notion of computable function pertains to what can be computed by the human brain, and (ii) that brains perform computations in the sense that is relevant to CTT. Premise (i) is questionable, but more importantly, premise (ii) is a consequence of McCulloch and Pitts's theory of the brain, which was not empirically plausible even when it was formulated.[3] So there is little reason to accept McCulloch and Pitts's justification of CTT. This, however, is no reason to abandon CTT, which remains well supported by the usual intuitive considerations.

McCulloch and Pitts's result – viz., that the neural networks defined within their theory, when supplemented with tapes, are computationally equivalent to TMs – was given a positive spin by John von Neumann. He reversed McCulloch and Pitts's order of justification, appealing to CTT in order to support a form of computationalism:

> The McCulloch–Pitts result … proves that anything that can be exhaustively and unambiguously described, anything that can be completely and unambiguously put into words, is ipso facto realizable by a suitable finite neural network. Since the converse statement is obvious, we can therefore say that there is no difference

between the possibility of describing a real or imagined mode of behavior completely and unambiguously in words, and the possibility of realizing it by a finite formal neural network. The two concepts are coextensive. A difficulty of principle embodying any mode of behavior in such a network can exist only if we are also unable to describe that behavior completely (von Neumann 1951, pp. 22–3).

Von Neumann's terminology, "describing behavior completely and unambiguously," is regrettably ambiguous, which makes it hard to know how von Neumann's argument is supposed to work. Under one possible reading, it is an extension of the legitimate argument given at the beginning: if neural activity is computation in the relevant sense, then, by CTT, there is a TM that computes it, and then, by McCulloch and Pitts's result, there is a McCulloch–Pitts network (possibly with tape) that computes it. This would be a sound argument, but it doesn't support the conclusion that neural activity is computation–rather, it presupposes it. Since von Neumann is clearly attempting to argue *for* computationalism ("any mode of behavior in such a network can exist"), he must mean something else.

Von Neumann's statement is perhaps the earliest example of what Jack Copeland has called the Church–Turing fallacy. This is the supposition that computationalism (or more weakly, the view that mental capacities can be simulated by TMs) follows from mechanism conjoined with CTT or some result established by Church and Turing (Copeland 1998, 2000; cf. also Kearns 1997).[4] CTT pertains to functions that are computable in the intuitive sense employed in mathematics. In order to show that something falls under CTT, it must first be shown that it is computable in that sense. CTT, per se, does nothing to establish that something is computable. Because of this, supposing that CTT entails computationalism is a fallacy. (The question of what can be *simulated*, in the sense of approximated, by TMs has nothing to do with CTT; it is addressed in Section 2 below.) And yet, as Copeland painstakingly shows, the fallacy is widespread in the cognitive science literature.

Here is an example by two philosophers:

[A] standard digital computer, given only the right program, a large enough memory and sufficient time, can compute any rule-governed input–output function. That is, it can display any systematic pattern of responses to the environment whatsoever (Churchland and Churchland 1990, p. 26).[5]

There are many problems with this passage; I will only discuss the most obvious one. The Churchlands appear to be saying that computers are universal, namely, they can compute any Turing-computable function (until they run out of memory and time). But by using the ambiguous (and vague) phrases "rule-governed functions" and "systematic patterns of responses," the Churchlands suggest something much stronger – and radically false. For if we accept CTT, we know from Turing's results and from standard computability theory that only countably many functions defined over strings, out of uncountably many, are computable by standard digital computers. And yet, in computability theory there is a clear sense in which many functions defined over strings that are not Turing-computable, such as the halting function, are "rule-governed" and "systematic". In this sense (among others), it is far from true that computers can compute all "rule-governed," "systematic" functions. The Churchlands's statement is a straightforward example of the kind of language that leads to the Church–Turing fallacy.[6]

If we want to make progress in the debate over computationalism, one thing we need is to eradicate the Church–Turing fallacy. And in order to eradicate it, it is not enough to expose it as fallacious. For several authors have attempted to go beyond von Neumann's quick remark and offer explicit arguments for computationalism based on CTT. These arguments need to be carefully evaluated – it takes some work to show why they fail.

Copeland contrasts computationalism with what might be called hypercomputationalism. According to hypercomputationalism, the brain computes non-Turing-computable functions. I wish to draw a more general contrast. I wish to contrast computationalism with the view that the brain does *not* compute Turing-computable functions. If the brain does not compute Turing-computable functions, this may be for either of two reasons. One is that, as hypercomputationalism holds, the brain computes non-Turing-computable functions. The other is that the brain does not compute anything at all – neural activity is something other than computation. The view that the brain does not compute Turing-computable-functions is still a form of wide mechanism in Copeland's sense, but it is more encompassing than Copeland's, because it includes both Copeland's hypercomputationalism and the view that mental capacities are not explained by neural computations but by neural processes that are not computational. Perhaps brains are simply not computing mechanisms but some other kinds of mechanisms. This view fits well with contemporary theoretical neuroscience, where much of the most rigorous and sophisticated work assigns no explanatory role to computation (cf. Dayan and Abbott 2001).

In order to assess CTT's relevance to computationalism, it is convenient to formulate computationalism in terms of Turing-computable functions. This can be done as follows:

(C) The functions from neural inputs to neural outputs are Turing-computable.

Given this reformulation of computationalism, CTT is potentially relevant to it, for CTT states that functions belonging to a certain class are Turing-computable. If the functions whose values are generated by brains belong to the relevant class of functions, then by CTT, (C) follows. There are three main ways in which serious arguments from CTT to computationalism have been run. Although they are old arguments, they keep reappearing in the literature without encountering adequate refutation. I address them in turn.

1. PHYSICAL CTT

An important and uncontroversial result of philosophical work on CTT during the last two or three decades is the distinction between CTT properly so called, which pertains to functions that are effectively calculable in the intuitive sense, and Physical CTT, which pertains to functions whose values are generated by physical systems. Whether or not they are aware of the distinction between CTT properly so called and Physical CTT, several authors have used Physical CTT to argue for (C). According to Physical CTT, all physically computable functions are Turing-computable. Since brains are physical systems, it follows from Physical CTT that the functions physically computed by brains are Turing-computable. So Physical CTT appears to entail (C).[7] In evaluating this argument from Physical CTT, we should distinguish between two importantly different — though seldom distinguished versions of Physical CTT. Accordingly, we need to examine two versions of the argument from Physical CTT.

1.1. *From Modest Physical CTT*

A first version of Physical CTT pertains to what functions can be *computed* by a mechanism or machine. Modest Physical CTT does not apply to all physical processes, but only to processes that are computations. It says that *the functions whose values are generated* *by computing mechanisms are Turing-computable*. In other words, if a mechanism performs computations, then Modest Physical CTT entails that that mechanism will compute functions that are Turing-computable.

Modest Physical CTT is relatively controversial. It is true if and only if genuine hypercomputers – machines that compute functions that are not Turing-computable – are physically impossible, and whether genuine hypercomputers are physically possible remains an open question.[8] Hypercomputers are an interesting theoretical possibility, and they are useful in discussions in the foundations of physics. Nevertheless, there is little if any evidence that genuine hypercomputers can be built and used by humans. In so far as it concerns computer scientists, Modest Physical CTT is quite plausible.

The version of Modest Physical CTT that concerns computer scientists is also the one that concerns neuroscientists and psychologists. For neuroscientists and psychologists are interested in neural and psychological mechanisms, which are relatively small physical systems confined within relatively small spatiotemporal regions. There is little reason to believe that neural mechanisms have access to the exotic physical resources, such as Malament–Hogarth spacetimes, that are exploited in designs for hypercomputers.

If Modest Physical CTT applies to brains, the resulting version of computationalism is not trivial. For Modest Physical CTT applies to the functions *computed* by physical systems, hence it entails that those systems are genuine computing mechanisms, whose activities are computations – as opposed to non-computing mechanisms, whose activities are not computations. In this respect, if we could conclude that Modest Physical CTT applies to brains, we would learn something substantive about them.

But Modest Physical CTT says nothing about whether any particular physical system is a computing mechanism. It leaves open whether the solar system, the weather, or your brain is a computing mechanism. Whether the brain or any other physical system is a computing mechanism must be established by means other than Modest Physical CTT. If we can establish that brains are computing mechanisms by other means, then Modest Physical CTT applies to them, and if Modest Physical CTT is true, then the functions brains compute are Turing-computable. What Modest Physical CTT establishes (if true) is only that if brains are computing mechanisms, then they are not hypercomputers.

1.2. *From Bold Physical CTT*

A second version of Physical CTT pertains to all physical systems, whether or not they perform computations. Bold Physical CTT says that *the functions whose values are generated (by computation or any other means) by physical systems are Turing-computable*. Hence, assuming that brains are physical, Bold Physical CTT does entail that the functions whose values are generated by brains are Turing-computable, which establishes (C). But this is a Pyrrhic victory.

To begin with, Bold Physical CTT is falsified by any genuinely random process. For as Turing knew well, genuinely random processes are not Turing-computable (cf. Piccinini 2003b).[9] Even if we restrict it to deterministic systems, Bold Physical CTT is difficult to make precise. The main reason is that the mathematical functions that are normally used to describe physical systems are functions of real (continuous) variables, whose domain and range include uncountably many values, whereas Turing-computable functions are functions of discrete variables, whose domain and range include only countably many values. Because of this, functions of real variables cannot be directly mapped onto Turing-computable functions.

There are several ways in which computability theory has been extended to functions of real variables. One such extension defines primitive computational operations that manipulate real-valued quantities instead of the strings of symbols of classical computability theory (Blum et al., 1998). Another proposal maintains the usual computations over strings of symbols but allows computations to rely on the exact values of real-valued constants (Siegelmann 1999). We may call the functions that are computable under these extensions of computability theory *real-computable* functions. Under these extensions, Bold Physical CTT may be reformulated as stating that a function is real-computable if and only if it is Turing-computable. Unfortunately, this version of Bold Physical CTT is far from true. Under either of the above extensions of computability theory, all functions from strings to strings – including all those that are not Turing-computable – are real-computable. So this formulation of Bold Physical CTT is false for rigorous mathematical reasons. The argument from Bold Physical CTT to (C) is valid but unsound.[10]

Even if there were a true version of Bold Physical CTT that could be used to entail (C) for some significant class of systems, however, this would yield only cold comfort to the computationalist. The main price of using Bold Physical CTT to support computationalism is that computationalism is thereby trivialized. The original motivation for computationalism is that the notion of computation can be used to *distinguish* mental processes from other processes – to find a mechanistic explanation that is specific to mental capacities (e.g., cf. Fodor 1998). But Bold Physical CTT cannot do this, because it applies indifferently to brains as well as other physical systems by virtue of their being physical. Any view of the brain derived from Bold Physical CTT is not a genuine form of computationalism, according to which mental capacities are *explained* by neural computations as opposed to some non-computational process. In theorizing about mental capacities, we are looking for mechanistic explanations that are specific to them. If computation is used in a sense that applies to any physical process, then it cannot be the basis for a specific explanation of mental capacities. So anyone who wishes to claim that brains are computing mechanisms in a sense that is specifically suited to explaining mental capacities, even if she believes Bold Physical CTT, must look for a more stringent version of computationalism and support it independently of Bold Physical CTT.

To summarize, neither version of Physical CTT helps the supporter of computationalism.

2. BETWEEN MODEST AND BOLD PHYSICAL CTT

In the study of physical systems and the properties of their mathematical descriptions, computability is relevant in several ways that do not fit comfortably within discussions of Physical CTT. I will briefly review some of them, which I think are more relevant to computationalism, and more fruitful to discuss, than Bold Physical CTT.

2.1. *Mathematical Tractability*

It may be useful to remind ourselves of the main reason why scientists, either in cognitive science or in any other sciences, resort to computational descriptions. Generally speaking, the need for computational descriptions in science has nothing to do with computationalism or the attempt to explain phenomena computationally. It has to do with the analytical intractability of most mathematical dynamical descriptions.

Given a system of equations describing a physical system, an analytic solution is a formula such that, given any initial condition of the system and any subsequent time *t*, the formula yields the state of the system at time *t*. The question of what systems can be solved analytically is not directly relevant to computability, but its answer leads to the important issue of computational approximation of physical systems, which is relevant to computability. It is well known that most systems of equations have no analytic solutions. In particular, the majority of nonlinear systems, which make up the majority of systems of equations, are not solvable analytically.[11]

In order to study systems that are not analytically solvable, a geometrical, qualitative approach has been developed by mathematicians.[12] This approach allows mathematicians to identify important qualitative features of a system's state space (e.g., its fixed points) without solving the system analytically. Unfortunately, the geometrical approach is suitable only for relatively simple systems, in which the number of state variables can be reduced to (at most) three, one per axis of a three-dimensional space. This limitation is mainly due to the fact that (ordinary) humans are unable to visualize a space with more than three-dimensions, and hence to apply this geometrical approach to systems whose state variables cannot be reduced to less than four. Nevertheless, the development of these geometrical techniques remains a fertile area of mathematical investigation. Overcoming the limitations of current methods for studying complex nonlinear dynamical systems, either by extending existing methods or by inventing new methods, is a current research project of many mathematicians.[13]

There is yet another way to tackle dynamical systems, whether solvable or unsolvable analytically, simple or complex. It is the use of computational methods for approximating the dynamical evolution of physical systems. The modern study of dynamical systems has exploded over the last half-century, to a large extent, thanks to the advent of digital computers. This is because computers, by offering larger and larger amounts of memory and computation speed, allow scientists to develop methods of approximation for systems of equations that are not analytically solvable, so as to study their behavior based on those approximation. The tool of computational approximation, which is one of the crucial tools in contemporary science, is what we now turn to.

2.2. *Computational Approximation*

There are at least two importantly different ways to approximate the behavior of a system computationally. One relies on the equations describing the dynamics of the system and on numerical methods for finding successive states of the system from those equations. The other does not rely on equations but treats the dynamics of the physical system itself as discrete. I will now briefly discuss these two methods.

Given a system of equations describing a physical system, whether or not the system is analytically solvable, it may be possible to develop methods for computing approximations of the behavior of the system. Working out specific numerical methods for specific sets of equations and showing that the resulting approximations are accurate within certain error bounds is another fertile area of mathematical investigation. These numerical methods, in turn, are behind the now widespread use of most computational models in science. These models are computer programs that exploit appropriate numerical methods to compute representations of subsequent states of a system on the basis of both the equations representing the system's dynamical evolution and data representing the system's initial conditions. Models of this kind can be constructed for any system whose behavior is described by known systems of equations that can be approximated by known numerical methods. This kind of computational approximation is perhaps the most popular form of contemporary scientific modeling.[14]

A different method of computational approximation relies on a computational formalism called cellular automata. Cellular automata are lattices of cells, each of which can take a finite number of discrete states and changes state in discrete time. At any given time step, the state of each cell is updated based on the state of its neighboring cells at that time. Different updating rules give rise to different cellular automata. In modeling a physical system using cellular automata, the system is spatially discretized in the sense that distinct spatial regions of the system are represented by distinct cells of a cellular automaton, and the dynamics of the system is temporally discretized in the sense that the state changes of the system's spatial regions are represented by the updating of the cells' states.

The pattern generated by the cellular automaton can then be compared with the observations of subsequent states of the system, so as to evaluate the accuracy of the approximation.[15]

The popularity and usefulness of computational approximations of physical systems, not only in physics but in many other sciences, may have been a motivating factor behind Bold Physical CTT. Some authors state forms of CTT according to which every physical system can be "simulated," by which they appear to mean computationally approximated in the present sense, by TMs.[16] But the question of whether every system can be computationally approximated is only superficially similar to Bold Physical CTT.

The importance of computational approximation is not that it embodies some thesis about physical systems and how to explain their behavior, but that it is the most flexible and powerful tool ever created for scientific modeling. An approximation may be closer or farther away from what it approximates. The way in which and the degree to which an approximation should mimic the system it models is largely a pragmatic factor, which depends on the goals of the investigators who are building the model.

If one allows computational approximations to be arbitrarily distant from the dynamical evolution of the system being approximated, then the thesis that every physical system can be computationally approximated becomes trivially true. If one is stricter about what approximations are acceptable, then that same thesis becomes nontrivial but much harder to evaluate. Formulating stricter criteria for acceptable approximations and evaluating what systems can be approximated to what degree of precision is a difficult question, which would be worthy of systematic investigation. Here, I can only make a few obvious points.

First, strictly speaking, unpredictable (e.g., non deterministic) systems cannot be computationally approximated. A computational approximation can only indicate the possible dynamical evolutions of such systems, without indicating which path will be followed by any given system.

Second, if there are any (deterministic or non deterministic) physical systems whose state transitions are not Turing-computable, e.g. if genuine hypercomputers are possible, then there is a strict sense in which those systems cannot be computationally approximated (by current computational methods).

Finally, as soon as the state-variables of a system are more than two and they interact nonlinearly in a sufficiently complex way, the system may exhibit chaos (in the mathematical sense). As is well known, chaotic systems are so sensitive to initial conditions that their dynamical evolution can only be computationally approximated for a relatively short time before diverging exponentially from the observed behavior of the system.

In conclusion, the extent to which physical systems can be computationally approximated depends both on the properties of physical systems and their mathematical descriptions, and on the criteria that are adopted for adequate approximation. The same computational model may count as producing adequate approximations for some modeling purposes but not for others. At any rate, on any nontrivial criteria for adequate approximation, it is far from true that every physical system can be computationally approximated. Having thus clarified the relation between computation and physical systems, we can go back to arguments from CTT to computationalism that do not rely on *Physical* CTT.

3. MENTAL PROCESSES AS THE FOLLOWING OF AN EFFECTIVE PROCEDURE

A straightforward way of arguing from CTT to (C) would be to show that mental processes are effective computations in the sense analyzed by Church and Turing. If mental capacities are effective in this sense, then the functions whose values are generated by brains when they exhibit mental capacities fall under CTT properly so called. Then, by CTT itself, (C) follows.[17] This section evaluates the thesis that mental processes are the following of an effective procedure.

Several authors believe that the claim that mental processes are the following of an effective procedure is an empirical hypothesis, to be supported on empirical grounds such as the successes of AI, psychology, or linguistics. This view does not concern us here. Here, I only discuss arguments that attempt to establish that mental processes are the following of an effective procedure without waiting for the sciences of mind and brain to run their course. The most explicit of such arguments is due to Judson Webb (1980, p. 236ff.; a similar argument is in Baum 2004, pp. 33–47).

Webb introduces his argument as if it were an explication of Turing's argument for CTT (ib., p. 220ff.). Webb is not alone in reading

Turing as offering an argument for computationalism,[18] but this is a misunderstanding of Turing. Turing expressed the following view:

> If the untrained infant's mind is to become an intelligent one, it must acquire both discipline and initiative. So far [i.e., by discussing effective procedures] we have been considering only discipline ... But discipline is certainly not enough in itself to produce intelligence. That which is required in addition we call initiative ... Our task is to discover the nature of this residue as it occurs in man, and try to copy it in machines (Turing 1948, p. 21).

There is no room here for a careful reconstruction of Turing's thought on intelligence and cognition. Suffice it to say that the consensus among Turing scholars is that in arguing for CTT, all that Turing was attempting to establish is that human *computation* processes are computable by TMs – he was not attempting to establish that all mental processes are the following of an effective procedure.[19]

Regardless of what Turing thought about this matter, here is what Webb says:

> To show that man is not an abstract (universal) Turing machine it would be sufficient to show that at least one of the following conditions is false:
>
> (i) Man is capable of only a finite number of internal (mental or physical) states $q_i \in Q$.
> (ii) Man is capable of discriminating only a finite number of external environmental states $s_i \in S$.
> (iii) Man's memory is described by a function f from pairs $\langle q_i, s_i \rangle$ to Q.
> (iv) There is a finite set B of atomic human behaviors, including some which may be identified with states s_i of S, which they effect (as in printing a symbol), and each $\langle q_i, s_i \rangle$ determines a unique element of B. A human (molar) behavior is comprised of a finite sequence of atomic behaviors of B, some of which have neural dimensions, while others may be molar behaviors in their own right (Webb 1980, p. 236).

Webb seems to believe conditions (i)–(iv) obtain. He should add that for his argument to go through, at least two further implicit conditions must obtain. First, the function from pairs $\langle q_i, s_i \rangle$ to B implicit in (iv) must be Turing-computable.[20] Second, humans must be capable of going through at most finitely many memory states and atomic behaviors in a finite time. It is well known that systems that do not satisfy these further conditions can compute non-Turing-computable functions (Giunti 1997; Pitowsky and Shagrir 2003). If Webb's explicit and implicit conditions are satisfied by human brains, then their behaviors are Turing-computable.

As to (i), Webb appeals to Turing's argument that a human who is performing calculations is capable of only finitely many internal states, otherwise some of them would be arbitrarily close and would be confused (Webb 1980, p. 221). This is justified in an analysis of human calculation, where the internal states must in principle be unambiguously identified by the computing humans on pain of confusion in performing the calculation. In his argument, Turing makes it clear that his "internal states" should be replaceable by explicit instructions.[21] Since the instructions have finite length, they can only distinguish between finitely many states.

But this says nothing about the number of internal states humans are capable of outside the context of calculation. Ordinary mathematical descriptions of physical systems ascribe to them uncountably many states. There is no *a priori* reason to suppose that humans are different from other physical systems in this respect. In fact, theoretical neuroscientists make extensive use of ordinary mathematical descriptions, which ascribe to neural mechanisms uncountably many states (Dayan and Abbott 2001).

As to (ii), Webb also attributes it to Turing. But again, Turing was only concerned with effective calculability by humans, an activity that must be describable by an effective procedure. Since effective procedures are finite, it seems plausible that they can only be used to discriminate between finitely many "environmental states" (i.e., symbols). Again, this says nothing about how many environmental states humans can discriminate outside the context of calculation.

Webb adds that (ii) "would follow from a finiteness condition in physics itself to the effect that there *were* only finitely many states of the environment there to be discriminated" (Webb 1980, p. 236). Webb, however, gives no reason to believe that such a physical condition obtains.

As to (iii), Webb makes it clear that the function f from pairs $\langle q_i, s_i \rangle$ to Q should be Turing-computable. He admits the possibility that f be nondeterministic, but submits that that could be taken care of by a nondeterministic TM (that is, a TM whose state transitions are not deterministic). But nondeterministic TMs can take care of this situation only if there are at most finitely many q_i and s_i, i.e. if (i) and (ii) obtains, and we've already seen that (i) and (ii) are unjustified. Webb is also skeptical that it could be "shown effectively" that there is no such Turing-computable f. If we assume that by "showing something effectively," Webb means proving something rigorously, Webb's statement is surprising. For one of Turing's greatest achievements was precisely to prove rigorously that there is no Turing-computable f

for solving uncountably many problems, including the halting problem for TMs. Given CTT, there is no principled difficulty in showing that a function is not Turing-computable.

The main difficulty with (iii), however, is not that *f* may not be Turing-computable, which there is no evidence for. The main difficulty is that Turing-computability may be simply irrelevant to "describing" human memory (in an explanatorily relevant way). The relationship between human memory mechanisms and Turing-computability can be divided into two sets of issues. One set of issues belongs with a general analysis of the relationship between computability and physical systems. This set of issues has nothing to do with whether mental processes in particular are the following of an effective procedure, and its relevance to computationalism was already covered in Sections 1 and 2. The other set of issues belongs with assessing the empirical hypothesis that the brain is a computing mechanism. That empirical hypothesis is not something that can be settled *a priori*.

As to (iv), Webb says, "it is … hard to imagine anything but Descartes' mechanical organisms existing at the dawn of evolution" (ib., 236), where the context makes clear that "Descartes' mechanical organisms" satisfy condition (iv). But the fact that something is "hard to imagine" is hardly conclusive evidence for (iv).

As to the further conditions implicit in Webb's argument, they seem no more *a priori* true than those explicitly stated by Webb. In the end, Webb has offered little support for his view that mental processes are the following of an effective procedure. This, of course, is not to say that mental processes are the following of a non-effective procedure, or a non-Turing-computable procedure. It may be that procedures are just irrelevant to scientific theories of mind and brain. To determine the relevance of procedures, whether effective or not, to neural or psychological theories, it seems more fruitful to develop and examine empirical theories of mind and brain rather than arguing *a priori* about these matters.

4. EFFECTIVE PROCEDURES AS A METHODOLOGICAL CONSTRAINT ON PSYCHOLOGICAL THEORIES

A final argument from CTT to (C) invokes a methodological constraint on psychological theories, to the effect that psychological theories should only be formulated in terms of effective procedures.

If this is the case, then by CTT, (C) follows. This argument is originally due to Daniel Dennett, and it has found many followers.[22] Here is Dennett's original:

[C]larity is ensured for anything expressible in a programming language of some level. Anything thus expressible is clear; what about the converse? Is anything clear thus expressible? The AI programmer believes it, but it is not something subject to proof; *it is, or it boils down to, some version of Church's Thesis* (e.g., anything computable is Turing-machine computable). But now we can see that the supposition that there might be a non-question-begging non-mechanistic psychology gets you nothing, unless accompanied by the supposition that Church's Thesis is false. For a non-question-begging psychology will be a psychology that makes no ultimate appeals to unexplained intelligence, and that condition can be reformulated as the condition that whatever functional parts a psychology breaks its subjects into, the smallest, or most fundamental, or least sophisticated parts *must not be supposed to perform tasks or follow procedures requiring intelligence*. That condition in turn is surely strong enough to ensure that *any procedure admissible as an "ultimate" procedure in a psychological theory falls well within the intuitive boundaries of the "computable" or "effective" as these terms are presumed to be used in Church's Thesis*. The intuitively computable functions mentioned in Church's Thesis are those that "any fool can do," while the admissible atomic functions of a psychological theory are those that "presuppose *no* intelligence." If Church's Thesis is correct, then the constraints on mechanism are no more severe than the constraints against begging the question in psychology, for any psychology that stipulated atomic tasks that were "too difficult" to fall under Church's Thesis would be a theory with undischarged homunculi [fn. omitted]. So our first premise, that AI is the study of all possible modes of intelligence, is supported as much as it could be, which is *not quite* total support, in two regards. The first premise depends on two unprovable but very reasonable assumptions: that Church's Thesis is true, and that there can be, in principle, an adequate and complete psychology (Dennett 1978a, p.83; emphasis added).

Dennett asserts that CTT yields a methodological constraint on the content of psychological theories. This is because, he says, any theory that postulates operations or procedures that are "too difficult" to fall under CTT is postulating an "undischarged homunculus," that is, an unexplained intelligent process. And any psychological theory that postulates undischarged homunculi should be rejected on the grounds that it begs the question of explaining intelligence.[23]

Dennett's reference to "procedures admissible as ultimate procedures in a psychological theory" implies that psychological theories are formulated in terms of procedures. Since what fall under CTT are *effective* procedures, Dennett's argument entails that the only ingredients of psychological theories are effective procedures. As a matter of fact, with the rise of cognitive psychology, some psychologists did propose that psychological theories be formulated as effective procedures, or computer programs, for executing the behavioral tasks explained by the theories (Miller et al., 1960). This view was elaborated philosophically by Fodor (1968), which is one of the works referred to by Dennett (1978a).

To the extent that psychological theories are or should be formulated in terms of effective procedures, Dennett's appeal to CTT is well motivated. For suppose that a psychologist offered an explanation of a behavior that appeals to a non-mechanical effective procedure of the kind imagined by Gödel (1965), or to arbitrary correct means of proof like those hypothesized by Kálmar (1959). Suppose that this psychologist refused to give effective instructions for these putative procedures. Then Dennett would be justified in concluding that these putative procedures are undischarged homunculi. Such a psychological theory purports to explain a behavior by postulating an unexplained intelligent process, which begins an infinite regress of homunculi within homunculi. Gödel and Kálmar's proposals may have a legitimate role to play in the philosophy of mathematics, but not in a naturalistic explanation of behavior, as a non-question-begging psychological explanation should be. In other words, any psychologist who wants to postulate effective procedures that do not fall under CTT should give a mechanistic explanation of how they can be followed by people. By so doing, this psychologist would falsify CTT. So far, Dennett's argument is sound.[24]

It remains to be seen the extent to which psychological theories are or should be formulated in terms of effective procedures. In Piccinini 2003a, I argue that explaining a behavior by postulating an effective procedure (or computer program) is only one species of a larger genus. This genus is explanation of behavior by mechanistic explanation, which consists of postulating a set of components and ascribing functions and organization to those components. In Piccinini (2004b), I also argue that in the philosophy of psychology tradition that goes from Fodor to Dennett and beyond, explanation by appeal to effective procedure and mechanistic explanation have been conflated. One of the effects of this conflation is Dennett's conclusion that CTT constitutes a methodological restriction on all psychological explanation rather than only on explanations that appeal to effective procedures (or computer programs).

Mechanistic explanations in psychology face the same constraint against begging the question of explaining intelligence that Dennett exploits in his argument. That is to say, the components postulated by a psychological mechanistic explanation should not contain undischarged homunculi. If a component is ascribed intelligence (or other high level cognitive abilities), this intelligence should be discharged by the lesser intelligence (or other high level cognitive abilities) of its components, until a level of analysis whose components have no intelligence (or other cognitive abilities) is reached. This methodological restriction on psychological mechanistic explanation was already formulated by Attneave (1961), who did not mention either effective procedures or CTT. He did not mention them because to the extent that psychological theories are formulated without postulating effective procedures, CTT is irrelevant to them.

CTT is only relevant to mechanistic explanations that postulate effective procedures, not to other kinds of mechanistic explanations. And *vice versa*: a psychological theory that explains behavior without postulating effective procedures does nothing by itself to falsify CTT. Modest Physical CTT, however, is relevant to any mechanistic explanation that postulates a process of computation. If a psychological theory postulates a genuine hypercomputation as a psychological process, then it falsifies Modest Physical CTT. If it postulates no computations at all, then it is irrelevant to Modest Physical CTT too. Finally, the issues of computational approximation discussed in Section 2 (Between Modest and Bold Physical CTT) are relevant to any psychological mechanistic explanation. They are relevant because they are relevant to any dynamical description – they have nothing in particular to do with psychological theories. In conclusion, CTT poses a methodological constraint on a species of psychological theories – those that postulate effective procedures – but poses no general constraint on psychological theories.

5. CONCLUSION

This paper addressed three attempts to support computationalism on the grounds of CTT. I argued that given a proper understanding of CTT, all those arguments are unsound. CTT does entail that *if* the brain follows an effective procedure, *then* that procedure is Turing-computable. And Modest Physical CTT does entail that *if* the brain performs computations, *then* those computations are Turing-computable. But neither CTT nor Modest Physical CTT is of any use in determining whether the brain follows effective procedures or more generally, whether it performs computations.

There is another way that computability is relevant to the explanation of mental capacities. Neural mechanisms are complex nonlinear dynamical systems *par excellence*, which lie at the frontier of what is mathematically analyzable by dynamical systems theory (Dayan and Abbott 2001; cf. Barabási 2002; Strogatz 2003). Because of this, methods of computational modeling are crucial to their scientific study. This is independent of whether neural mechanisms are computing mechanisms in any nontrivial sense.

Where does this leave computationalism? Computationalism is one family of theories of mind and brain among others. The science of mind and brain belongs with physiology and engineering, which explain the behavior of systems by finding mechanisms (Machamer et al., 2000). Some mechanisms perform computations (e.g., digital computers) and some don't (e.g., stomachs). Computationalism is true if and only if neural mechanisms perform computations and those computations explain mental capacities. The only way to find out is to investigate the properties of neural mechanisms empirically and search for explanations of mental capacities.

ACKNOWLEDGMENTS

Thanks to Darren Abramson and two referees for their helpful comments.

NOTES

[1] Computationalism is a form of physicalism (or mechanism, as it is often referred to in the literature on computationalism). Different forms of physicalism may differ on which mental processes are identical to or realized by physical processes and whether the realizing processes are narrow (inside the organism) or broad (extending into the environment). Furthermore, different versions of computationalism may differ on which mental capacities are explained by inner computations. For present purposes, nothing hinges on these differences.

[2] I am appealing to the standard understanding of CTT, according to which CTT provides a correct analysis of the intuitive notion of computability employed by mathematicians. There is no room here to summarize the vast literature on CTT. Useful references include Kleene (1952), Copeland (2002a), Folina (1998), and Sieg (2001).

[3] For a detailed study of McCulloch and Pitts's theory, see Piccinini (2004a).

[4] The tendency to commit the Church-Turing fallacy is probably promoted, at least in part, by the common conflation between CTT and Turing's discovery that there are universal TMs, namely, TMs that compute any function computable by

ordinary TMs (for a recent example of such a conflation, see Wolfram 2002, p. 1125).

[5] Cf. also Guttenplan (1994, p. 595).

[6] For a similar dismissal of Churchland and Churchland's statement, see Copeland (2000, 2002a).

[7] Cf:

Human beings are products of nature. They are finite systems whose behavioral responses to environmental stimuli are produced by the mechanical operation of natural forces. Thus, according to Church's thesis, human behavior ought to be simulable by Turing machine (McGee 1991, p. 118).

[W]e have good reason to believe that the laws of physics are computable, so that we at least ought to be able to *simulate* human behavior computationally (Chalmers 1996, p. 313).

The use of the term "simulate" in these quotes may suggest that these authors are arguing for something weaker than computationalism. Perhaps they are suggesting that computations can generate approximations of mental capacities, without explaining them. But the term "computational simulation" is used to mean either computational approximation (as when a weather forecasting program simulates the weather) or computational replication (as when a universal TM simulates another TM). Unfortunately, these authors do not state explicitly which sense of "simulation" they are employing. Since they appeal to CTT, however, and since CTT is only relevant to the stronger notion of simulation (see Section 2 for more on computational approximation), they commit themselves to the stronger notion of simulation. At any rate, the context of their remarks makes clear that these authors do intend to argue for computationalism.

[8] Two recent reviews of the literature are Copeland (2002b) and Cotogno (2003).

[9] My point is simply that we cannot pick a Turing Machine (or any other process) and ask it to generate the same outputs as a genuine random process unless we already know all the outputs of the random process. (Even knowing all the outputs may not help, because most of the sequences generated by genuinely random processes are not Turing-computable).

[10] Another way to link computation to real-valued processes asks whether, when a deterministic system's initial conditions are defined by a computable real number (i.e., a real number that can be printed out by a TM), the system's dynamical evolution always leads to states defined by computable real numbers. This question is quite removed from the usual concerns of computationalism, but at any rate, there are field equations for which the answer is negative (Pour-El and Richards 1989).

[11] Sometimes, the question of whether a system of equations is analytically solvable gets confused with whether a function is computable:

Most dynamical systems found in nature cannot be characterized by equations that specify a computable function. Even three bodies moving in Newtonian space do not satisfy this assumption. It is very much an open question whether the processes in the brain that subserve cognition can be characterized as the computation of a computable function (Cummins 2000, p. 130).

This brief argument is a nice counterpoint to the arguments by McGee and Chalmers cited above. Like Chalmers and McGee, Cummins is mistakenly assuming that the notion of computability applies directly to the functions described by ordinary dynamical descriptions, such as differential equations. Unlike them, Cummins is assuming that whether a system of equations is solvable is the same question as whether a function is computable. What Cummins should have said is that most dynamical systems are not characterized by equations that are analytically solvable.

[12] For an introduction, see Strogatz (1994).

[13] Two popular books that introduce some recent, groundbreaking work in this area are Barabási (2002) and Strogatz (2003).

[14] For more on computational models in science, see Humphreys (1990) and Rohrlich (1990).

[15] For more on cellular automata as a modeling tool, see Rohrlich (1990) and Hughes (1999).

[16] This is one possible interpretation of von Neumann's remark, cited above. Here are two recent examples: "Church's thesis [states] that anything that can be given a precise enough characterization as a set of steps can be simulated on a digital computer" (Searle 1992, p. 200); "theories of cognition are formulated in terms on processes that could be emulated by programs running on a digital computer" (Scott 1997, p. 68); see also Baum 2004, p. 47. The reason given by Scott for his view is that *all* scientific theories "must be the embodiment of a Turing computable function" (Scott 1997, p. 63), and his reason for that is that due to CTT together with the requirement that scientific theories be publicly understandable, "any genuinely scientific theory must embody an effective procedure that will allow its intended audience to determine what it entails about the range of situations to which it applies" (ibid., p. 66). Scott's argument is a non sequitur. For even if, for the sake of the argument, we accept the otherwise dubious premise that any scientific theory must embody an effective procedure for deriving its logical consequences, it doesn't follow that the processes *described* by the theory either are computational or may be simulated computationally.

[17] Cf: "Human cognitive processes are effective; by the [Church–Turing] thesis it follows they are recursive relations" (Nelson 1987, p. 581).

[18] Cleland (1993, p. 284), Shanker (1995, p. 55), Fodor (1998) and Baum (2004, p. 33).

[19] See Sieg (1994) and Copeland (2000). I have argued at length against this misinterpretation of Turing in Piccinini (2003b).

[20] This condition is entailed by Webb's conditions (i), (ii), and (iv), but I will soon question them.

[21] This point has been emphasized by Sieg (2001).

[22] Webb (1980, p. 220), Haugeland (1981, p. 2), Fodor (1981, pp. 13–15), Pylyshyn (1984, p. 52, 109) and Boden (1988, p. 259). A previous discussion of this argument can be found in Tamburrini (1997).

[23] For more by Dennett on homunculi and question begging in psychology, cf. Dennett (1978c, pp. 57–9, d, 119ff). For more by Dennett on AI as the study of all mechanisms (as opposed to computing mechanisms) for intelligence, see Dennett (1978b, esp. p. 112).

[24] Except for his implicit claim that a psychology that falsifies CTT would be a "non-mechanistic psychology". Unlike me, Dennett is implicitly operating with a narrow notion of mechanism, according to which only computational processes that are Turing-computable count as mechanistic. Needless to say, Dennett does not offer any defense of his implicit restrictions on what counts as mechanistic.

REFERENCES

Attneave, F.: 1961, 'In Defense of Homunculi', in W. Rosenblith (ed.), *Sensory Communication*, MIT Press, Cambridge, MA, pp. 777–782.

Barabási, A.-L.: 2002, *Linked: The New Science of Networks*, Cambridge, MA, Perseus.

Baum, E. B.: 2004, *What is Thought?*, Cambridge, MA, MIT Press.

Blum, L. and F. Cucker, et al.: 1998, *Complexity and Real Computation*, New York, Springer.

Boden, M. A.: 1988, *Computer Models of Mind: Computational Approaches in Theoretical Psychology*, Cambridge University Press, Cambridge.

Chalmers, D. J.: 1996, *The Conscious Mind: In Search of a Fundamental Theory*, Oxford University Press, Oxford.

Church, A.: 1936, 'An Unsolvable Problem in Elementary Number Theory', *The American Journal of Mathematics* **58**, 345–363.

Churchland, P. M. and P. S. Churchland: 1990, 'Could a Machine Think?' *Scientific American* **CCLXII**, 26–31.

Cleland, C. E.: 1993, 'Is the Church–Turing Thesis True?', *Minds and Machines* **3**, 283–312.

Copeland, B. J.: 1998, 'Turing's O-machines, Searle, Penrose, and the Brain', *Analysis* **58**(2), 128–138.

Copeland, B. J.: 2000, 'Narrow Versus Wide Mechanism: Including a Re-Examination of Turing's Views on the Mind-Machine Issue', *The Journal of Philosophy* **XCVI**(1), 5–32.

Copeland, B. J.: 2002a, 'The Church–Turing Thesis', *The Stanford Encyclopedia of Philosophy (Fall 2002 Edition)*. E. N. Zalta, URL = ⟨http://plato.stanford.edu/archives/fall2002/entries/church–turing/⟩.

Copeland, B. J.: 2002b, 'Hypercomputation', *Minds and Machines* **12**, 461–502.

Cotogno, P.: 2003, 'Hypercomputation and the Physical Church–Turing Thesis', *British Journal for the Philosophy of Science* **54**, 181–223.

Dayan, P. and L. F. Abbott: 2001, *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*, MIT Press, Cambridge, MA.

Dennett, D. C.: 1978a, *Brainstorms*, MIT Press, Cambridge, MA.

Dennett, D. C.: 1978b, 'Introduction', in D. C. Dennett (ed.), *Brainstorms*, MIT Press, Cambridge, MA, pp. xi–xxii.

Dennett, D. C.: 1978c, 'Skinner Skinned', in D. C. Dennett (ed.), *Brainstorms*, MIT Press, Cambridge, MA, pp. 53–70.

Dennett, D. C.: 1978d, 'Artificial Intelligence as Philosophy and as Psychology', in D. C. Dennett (ed.), *Brainstorms*, MIT Press, Cambridge, MA, pp. 109–126.

Fodor, J. A.: 1968, 'The Appeal to Tacit Knowledge in Psychological Explanation', *Journal of Philosophy* **65**, 627–640.

Fodor, J. A.: 1981, *Representations*, MIT Press, Cambridge, MA.

Fodor, J. A.: 1998, *Concepts*, Clarendon Press, Oxford.

Folina, J.: 1998, 'Church's Thesis: Prelude to a Proof', *Philosophia Mathematica* **6**, 302–323.

Giunti, M.: 1997, *Computation, Dynamics, and Cognition*, Oxford University Press, New York.

Gödel, K.: 1965, 'Postscriptum', in M. Davis. (ed.), *The Undecidable*, Raven, New York, pp. 71–73.

Guttenplan, S.: ed. 1994, *A Companion to the Philosophy of Mind*, Blackwell, Cambridge, MA.

Haugeland, J.: 1981, 'Analog and Analog', *Philosophical Topics* **12**, 213–225.

Hughes, R. I. G.: 1999, 'The Ising Model, Computer Simulation, and Universal Physics', in M. S. Morgan and M. Morrison (eds.), *Models as Mediators*, Cambridge University Press, Cambridge, pp. 97–145.

Humphreys, P.: 1990, 'Computer Simulations', *PSA 1990* **2**, 497–506.

Kálmar, L.: 1959, 'An Argument Against the Plausibility of Church's Thesis', in A. Heyting (ed.), *Constructivity in Mathematics*, North-Holland, Amsterdam. pp. 72–80.

Kearns, J.T. (1997). "Thinking Machines: Some Fundamental Confusions." *Minds and Machines* **7**: 269–287.

Kleene, S. C.: 1952, *Introduction to Metamathematics*, Princeton, Van Nostrand.

Machamer, P. K., L. Darden and C. Craver: 2000, 'Thinking About Mechanisms', *Philosophy of Science* **67**, 1–25.

McCulloch, W. S. and W. H. Pitts: 1943, 'A Logical Calculus of the Ideas Immanent in Nervous Activity', *Bulletin of Mathematical Biophysics* **7**, 115–133.

McGee, V.: 1991, 'We Turing Machines Aren't Expected-Utility Maximizers (Even Ideally'), *Philosophical Studies* **64**, 115–123.

Miller, G. A., E. H. Galanter and K. Pribram: 1960, *Plans and the Structure of Behavior*, Holt, New York.

Nelson, R. J.: 1987, 'Church's Thesis and Cognitive Science', *Notre Dame Journal of Formal Logic* **28**(4), 581–614.

Piccinini, G. (2003a). Computations and Computers in the Sciences of Mind and Brain. Pittsburgh, PA, University of Pittsburgh. URL=<http://etd.library.pitt.edu/ETD/available/etd-08132003-155121/>

Piccinini, G. (2003b). "Alan Turing and the Mathematical Objection." *Minds and Machines* **13**(1): 23–48.

Piccinini, G. (2004a). "The First Computational Theory of Mind and Brain: A Close Look at McCulloch and Pitts's 'Logical Calculus of Ideas Immanent in Nervous Activity'." *Synthese* **141**(2): 175–215.

Piccinini, G. (2004b). "Functionalism, Computationalism, and Mental States." *Studies in the History and Philosophy of Science* **35**(4): 811–833.

Pitowski, I. and O. Shagrir: 2003, 'Physical Hypercomputation and the Church–Turing Thesis', *Minds and Machines* **13**(1), 87–101.

Pour-El, M. B. and J. I. Richards: 1989, *Computability in Analysis and Physics*, Springer Verlag, Berlin.

Pylyshyn, Z. W.: 1984, *Computation and Cognition*, MIT Press, Cambridge, MA.

Rohrlich, F.: 1990, 'Computer Simulation in the Physical Sciences', *PSA 1990* **2**, 507–518.

Scott, P. D. (1997). Crisis? What Crisis? Church's Thesis and the Scope of Cognitive Science. *Two Sciences of Mind: Readings in Cognitive Science and Consciousness.* S. Ó Nualláin, P. Mc Kevitt and E. Mac Aogáin. Philadelphia, John Benjamins: 63–76.

Shanker, S. G.: 1995, 'Turing and the Origins of AI', *Philosophia Mathematica* **3**, 52–85.

Sieg, W.: 1994, 'Mechanical Procedures and Mathematical Experience', in G. Alexander (ed.), *Mathematics and Mind*, Oxford University Press, New York, pp. 71–117.

Sieg, W.: 2001, 'Calculations by Man and Machine: Conceptual Analysis', in W. Sieg, R. Sommer and C. Talcott (eds.), *Reflections on the Foundations of Mathematics (Essays in Honor of Solomon Feferman)*, Association for Symbolic Logic, vol. **15**, pp. 387–406.

Siegelmann, H. T.: 1999, *Neural Networks and Analog Computation: Beyond the Turing Limit*, Birkhäuser, Boston, MA.

Searle, J. R.: 1992, *The Rediscovery of the Mind*, MIT Press, Cambridge, MA.

Strogatz, S. H.: 1994, *Nonlinear Dynamics and Chaos*, Perseus, Cambridge, MA.

Strogatz, S. H.: 2003, *Sync: The Emerging Science of Spontaneous Order*, Hyperion, New York.

Tamburrini, G.: 1997, 'Mechanistic Theories in Cognitive Science: The Import of Turing's Thesis', in M. L. Dalla Chiara and et al. (eds.), *Logic and Scientific Method*, Kluwer, pp. 239–257.

Turing, A. M.: 1936-7 [1965], 'On Computable Numbers, with an Application to the Entscheidungsproblem', in M. Davis. (ed.), *The Undecidable*, Raven, Ewlett.

Turing, A. M.: 1948, 'Intelligent Machinery', in D. Ince (ed.), *Mechanical Intelligence*, North-Holland, Amsterdam, pp. 87–106.

von Neumann, J.: 1951, 'The General and Logical Theory of Automata', in L. A. Jeffress (ed.), *Cerebral Mechanisms in Behavior*, Wiley, New York, pp. 1–41.

Webb, J. C.: 1980, *Mechanism, Mentalism, and Metamathematics*, Reidel, Dordrecht.

Wolfram, S.: 2002, *A New Kind of Science*, Wolfram Media, Champaign, IL.

Department of Philosophy
University of Missouri – St. Louis
599 Lucas Hall (MC 73)
One University Blvd.
St. Louis, MO 63121
USA
E-mail: piccininig@umsl.edu