

NAGW-1592
IN-61-CR
198504
14P.

Theoretical Note

Consequences of nonclassical measurement for the algorithmic description of continuous dynamical systems

Chris Fields

Knowledge Systems Group, Computing Research Laboratory
New Mexico State University, Las Cruces, NM 88003-0001, USA

Running head: Algorithmic description of dynamical systems

(NASA-CR-184851) CONSEQUENCES OF
NONCLASSICAL MEASUREMENT FOR THE ALGORITHMIC
DESCRIPTION OF CONTINUOUS DYNAMICAL SYSTEMS
(New Mexico State Univ.) 14 P CSCI 09B

N89-20640

Unclas
63/61 0198624

Abstract: Continuous dynamical systems intuitively seem capable of more complex behavior than discrete systems. If analyzed in the framework of the traditional theory of computation, a continuous dynamical system with countably many quasistable states has at least the computational power of a universal Turing machine. Such an analysis assumes, however, the classical notion of measurement. If measurement is viewed nonclassically, a continuous dynamical system cannot, even in principle, exhibit behavior that cannot be simulated by a universal Turing machine.

Introduction.

The recent resurgence of interest in the computational properties of both artificial and biological neural networks has been accompanied by a renewal of interest in the scope of the notion of computation as discrete manipulation that has its formal realization in the Turing machine, and in the question of whether this notion encompasses all of the functions that are computable by neural networks. A number of authors, in the course of emphasizing the qualitative and architectural differences between computations as carried out by artificial neural networks and as carried out by Turing or von Neumann machines have hinted, at any rate, that machines of the two types might compute different classes of functions (e.g. Kohonen, 1988, Ch. 9; Smolensky, 1988; Freeman, 1988). It is unclear whether these authors intend to question the Church-Turing thesis, or merely to make a point about the appropriateness of different architectures for different classes of functions. The stronger reading is certainly the more interesting, as the Church-Turing thesis is the current operational definition of the intuitive notion of computability. Were neural networks shown to be capable of computing functions that could not be computed by Turing or von Neumann machines, it is standardly assumed, a change would be required not only in the formalism of computability theory, but in the underlying intuitions as well.

Artificial neural networks are typically represented formally as dynamical systems that are discrete in space, but continuous in time (reviewed by Grossberg (1988); see also, e.g. Cohen (1984) for a discussion of "networks" that are continuous in space as well). The temporal continuity of computations by such networks is typically the focus of discussions of possible differences between their computational powers and those of traditional computers: continuous systems intuitively appear to possess more behavioral degrees of freedom, and hence possibly greater computational power, than do discrete systems. Smolensky (1988), for example, suggests that "if there is any new theory of computation implicit in the subsymbolic approach [to cognitive science], it is likely to be a result of a fundamentally different, continuous formulation of computation" (p. 18). Implicit in this suggestion is the idea that continuous dynamical systems may be capable of computations that are not representable as sequences of discrete manipulations, and hence not describable algorithmically.

The purpose of this note is to consider the description of continuous dynamical systems in light of the theory of measurement. While dynamical system models are often analyzed simply as formal devices, the more relevant issue in computer science or artificial intelligence is whether the observed behavior of a finite, physically-realized dynamical system can be described algorithmically. In practice, the process of measuring the states of any physically-realized dynamical system will, as pointed out by Ashby (1956), impose a discrete description on its behavior, for only finitely many states of the system can actually be observed. If the measurement process is regarded as classical, however, continuously many distinct discrete descriptions, each of which may approach countable complexity, are possible in principle. Because of this explosion of distinct possible descriptions, even relatively simple dynamical systems appear to have at least the computational power of a universal Turing machine. This observation appears to be the main motivation behind current questions regarding the Church-Turing thesis. If the measurements in question are regarded as nonclassical, however, this situation does not arise. A nonclassical view of measurement allows the coherent, algorithmic description of continuous dynamical systems both in the case in which the resolution of the measurements is sufficiently low that the behavior of the system can be described classically, and in the case of nonclassical system behavior.

Measurement in classical state spaces.

The *state space* associated with a Turing machine is the set of all possible configurations of both the head and the tape¹; in the case of a binary Turing machine with an infinite tape, the state space is isomorphic to the cartesian product of the set of all binary vectors of length L with the positive integers. The execution trace of the machine for a given input is a path in the state space, i.e. a finite sequence of head positions and tape configurations. The initial state of such a path is the input, which in the case of a universal Turing machine may specify both program and data, and the final state is the output.

In the case of an arbitrary virtual machine M , the state space is the set of all states accessible to the machine. The identities of these states are stipulated, at least implicitly, by the formal definition of M . An execution trace of M for a given input is a path in this space. Consider any such trace t , with an initial state t_i and a final state t_f that is one of the possible final

states of the machine. A procedure that generates this trace from t_i can be specified trivially:

if state = t_f , halt.

if state = t_j , $j \neq f$, state $\leftarrow t_{j+1}$.

Provided the execution traces are known, it is, therefore, trivial to simulate any virtual machine M with a universal Turing machine.

If a virtual machine M is to be realized in a physical system S, it must be possible to construct an isomorphism from the state space of M to a set of states of S that interprets all execution traces of M as sequences of states of S that are dynamically allowed, and that end in quasistable states, i.e. sequences s_i, \dots, s_f of states of S such that, if S is placed in s_i , it will traverse the states in the sequence until it reaches s_f , in which it will remain until further perturbed. The image of such an isomorphism in S is a representation of the computation carried out by M, i.e. an interpretation of S as realizing M. The construction of such an isomorphism requires, clearly, that suitable states of S have been identified. The identification of states of S that can serve as targets of a virtual machine interpretation is stipulative, but is *not arbitrary*: the states so identified must be measurable - reproducibly distinguishable from other states of S by observation.

A state s of a physical system S is *measurable* if there is another physical system S' with a state s' such that, if S and S' are allowed to interact, S' will be in state s' only if S is in state s . This definition is implicit in both classical physics and cybernetics. It is clearly regressive; the regression is halted, in practice, by reliance on unanalyzed instances of measurement by human observers. Thus the formal notions of computation and of measurement both depend, for their foundations, on the intuitive sense of a human observer that a physical system has reproducibly changed its state.

A system S is *classical* if its states can be measured without being perturbed by its interaction with the measuring system. If a system is classical, in other words, the process of measuring a given state has no effect on subsequent state transitions. Computers are designed, as are most artifacts, to behave as classical devices; it is possible, for example, to observe

representations of values of variables printed on a terminal during execution of a program without altering the flow of control of the program. Classical measurability is, moreover, assumed implicitly in the definition of a Turing machine, in that it is assumed that the states traversed by the machine in the course of a computation can be followed by an observer without the observations perturbing the machine in any way. The theory of computation thus assumes, not merely some notion of measurement, but the *classical* notion of measurement. This assumption of classical measurability renders the algorithmic description of continuous dynamical systems both seemingly straightforward and deeply paradoxical.

If a system is both classical and continuous, the precision with which its behavior over time can be described, i.e. the number of measurable states that the system can be interpreted as traversing in the course of exhibiting a particular behavior, is limited only by the resolution of the measuring system that is used to identify its states. The interpretation of a continuous, classical dynamical system S as a virtual machine M is, therefore, quite straightforward. For each execution trace t of M , one selects a sequence s_i, \dots, s_f of states of S that can be mapped one-to-one onto t . The only requirements of this interpretation process are that the terminal states s_f be quasistable, and that S have a sufficient number of distinct states. Provided that measurements can be made at sufficient resolution, however, a sufficient number of states can be identified in a classical system that is changing at least one of its state variables - e.g. the position of one of its components - continuously to allow any interpretation whatever. The behavior of a continuous, classical dynamical system between two states s_i and s_f can thus be interpreted as a trace of any computation by any virtual machine whatever.

Consider a continuous, classical dynamical system S with countably many quasistable states, each of which may be reached by a distinct path s_i, \dots, s_f . Each of these paths may independently be interpreted as a trace of a computation by some virtual machine. In particular, a path in S can be found to interpret every possible path in any finite virtual machine M . Such a system S can, therefore, be interpreted as executing any algorithm whatever on any input whatever, merely by measuring a sufficient number of paths at sufficient resolution. A continuous dynamical system with countably many quasistable states has, therefore, at least the computational power of a universal Turing machine. Moreover, the number of sets of measurements, and hence the number of discrete sets of states of a continuous system, that can be

given any particular virtual-machine interpretation is continuous.

The above result, which is an analog of Ashby's (1956, p. 39) point that any physical object can be given countably many interpretations as a "system," renders the question "what algorithm does S execute?" answerable only by stipulation. By explicitly interposing an interpretation step between the system's dynamics and its purely stipulative, explosively numerous algorithmic descriptions, it also permits the question of whether the dynamics underlying the interpretations can, in some sense independently of any interpretation, be regarded as having computational power. The intuition that continuous systems are more behaviorally complex than discrete systems suggests that the underlying dynamics have computational power that cannot be described algorithmically, and the current wave of questions about the algorithmic description of dynamical systems appears to be motivated by the consideration of particular cases of systems that exhibit very complex dynamic behavior. Smolensky (1988), for example, portrays algorithmically describable "symbolic" cognition as an emergent property of a complex underlying "subsymbolic" dynamical system in which temporal continuity plays a key role. The significance assigned to continuity suggests that the underlying system is viewed as not describable algorithmically. Similarly, Freeman (1988) explicitly contrasts algorithmic systems and chaotic dynamical systems, with the implication that the latter are not describable algorithmically.

The paradox imposed on computational description by the classical notion of measurement is that the "underlying dynamics" of a continuous system appear both accessible and forever inaccessible to the analyst. The behavior of a classical continuous system can be simulated exactly - even if its dynamics are everywhere nondifferentiable - at any level of resolution at which its states can be measured. Classical measurement can, however, be made arbitrarily precise; hence *any* behavior, no matter how complex, can be measured, and then simulated with a Turing machine. On the other hand, because the system is continuous, any particular discrete simulation potentially misses behavior of arbitrary complexity altogether. This paradox is particularly striking in the case of everywhere nondifferentiable systems, in which the analyst is driven to infinitely greater levels of resolution in an attempt to "capture" the behavior in an algorithmic description, but is forever frustrated. Ashby (1956), perhaps sensing this paradox, simply refused to consider such systems as "machine-like" and worthy of

analysis (p. 41).

Nonclassical measurement.

In practice, measuring systems have finite resolution, and systems being measured exhibit thermal and other forms of noise that further limit the precision with which states can be identified. Hence the dynamics of a continuous system can, in practice, only be sampled over a finite interval, and usually in a way that averages many states that would be distinct if viewed at higher resolution. These features of measurement in practice became features of measurement in principle in the course of the transition from classical to nonclassical physics, during which it was recognized that the classical notion of measurement, while useful as an approximation, is deeply incoherent. A system is *nonclassical* if the interactions with an external system required to measure its states perturb its dynamics, i.e. if some version of Heisenberg's principle holds².

Physical systems become nonclassical at the level of measurement resolution at which the interaction energy between the system and the measuring system becomes comparable with the separation in energy between the states being measured (a standard treatment of nonclassical measurement is Jauch, 1968, Ch. 11; more recent reviews may be found in Greenberger, 1986). As the energy transferred in a measurement process determines the information that may be transferred, more informative - i.e. more precise - measurements cause greater perturbations of the measured system. Note that "measurement" here does not imply an *epistemic* act by an intelligent agent, but merely an interaction with an external system that transfers energy - i.e. information - to the system being measured: any such interaction counts as a measurement. At measurement resolutions that are sufficiently low that the measurement process does not perturb the system, the behavior of the system can, as an approximation, be regarded as classical. While the theory of nonclassical measurement has been formulated mathematically as part of physics, it is generally accepted that biological and psychological systems, for example, are nonclassical with respect to many kinds of measurements. Indeed most measurements made on living systems perturb them in some way, and such perturbations cannot be predicted exactly in the absence of nonperturbing measurements. The nonclassical nature of such systems is rooted in the nonclassical nature of their physical components; hence increasing the

measurement resolution magnifies instead of mitigating their nonclassical behavior.

The identification of a state requires measurement; therefore, identifying states of a nonclassical system may perturb the system, and hence may cause a state transition to occur that would not have occurred had the state not been identified. In a nonclassical "Turing machine," for example, observation of the position of the head, or of the state of the tape, may influence the head's next move, or the next state of the tape. Neither the direction nor the magnitude of such perturbations can be predicted exactly, and the measurement of such a perturbation will only further perturb the system. As the resolution with which states are measured increases past the upper limit for classical behavior, each measurement introduces further uncertainty regarding the state transitions that the system will execute - i.e. the path that it will follow - from any given initial state. Higher resolution measurements require greater energy transfer, which causes greater perturbation of the dynamics, and hence reduces the predictability of future behavior. In the limit of an infinitely precise measurement of the state, all information about the next state transition is lost: all transitions become equally probable. The behavior of a system in the nonclassical domain can, therefore, only be specified stochastically, i.e. by specifying a probability distribution over the possible transitions from each state. That such stochastic specifications are *complete* descriptions, in principle, of the system's *observable* dynamics is a foundational, and experimentally well confirmed, claim of quantum mechanics (see e.g. Greenberger, 1986, esp. Parts III, IV, IX).

The treatment of measurements of the behavior of continuous dynamical systems as nonclassical has two effects. First, it places an upper limit on the resolution with which states can be measured, and hence the precision with which they can be specified, if classical behavior is to be maintained. By preventing the measurement resolution from approaching infinity in principle, such an upper limit prevents the number of states that can be identified as belonging to a given path from approaching infinity in principle. The explosion of interpretations that results from the assumption of classical measurement is thus prevented: because only a finite number of states can be identified on any path, there will always be some virtual machine that a given system does not have sufficient measurable states to realize. No physical system that behaves classically can, therefore, realize every virtual machine, even in principle. A *continuous* physical system that is measured in such a way that its behavior remains classical has, in other

words, fewer measurable states, and hence lesser computational power, than a *discrete* universal Turing machine. A universal Turing machine can, therefore, simulate the behavior of any continuous physical system, at any level of resolution at which the behavior of the latter can be regarded as classical. This is a consequence not of a fact about continuous dynamics *per se*, but of the requirements of nonclassical measurement.

The second consequence of the nonclassical treatment of measurement is that the behavior of a continuous dynamical system can be completely described algorithmically, even in the range of resolution in which the behavior is nonclassical. This follows from the completeness of the description of the system's dynamics by the state transition probability distributions. A "transition virtual machine" specified entirely in terms of such probabilities was described by Fields and Dietrich (1987) as a model of cross-domain problem solving behavior; such a virtual machine can be viewed intuitively as a stochastic ensemble of Turing machines, and can be simulated on a von Neumann machine by simply calculating the probability distributions.

The nonclassical view of continuous dynamical systems addresses the paradox posed by their classical description by affirming that only observed behavior can be described, and that only stochastic behavior can be observed. The description of a continuous dynamical system as a transition virtual machine is, therefore, a complete description of its observable dynamics at the level of resolution of the observations. As the level of resolution increases, the information about the dynamics contained in the description decreases; in the limit of infinite resolution, the system's behavior appears random, and thus contains no information. As measurement resolution decreases to below the classical limit, the behavior of the system appears progressively more determined. However, the number of identified states available for computational interpretation concomitantly decreases; the computational power of the system thus decreases with the resolution of the measurements. The "underlying dynamics" of the system cease to hover always on the edge of observability: there are no underlying dynamics, only observable - and hence algorithmically describable - dynamics.

Conclusions.

How one approaches the question of what is computable by a physical system depends on how one approaches the process of measurement. The standard Turing machine model of computation implicitly assumes that measurement is a classical process, i.e. that the states of a system can be measured arbitrarily precisely without perturbing the system. When applied to continuous dynamical systems, however, this view of measurement leads, by allowing the resolution with which states are identified to approach infinity, to the paradox that even relatively simple systems may be interpreted as computing any function, i.e. as having at least the computational power of a universal Turing machine.

Classical measurement is, however, an approximation that is valid only at low resolution. If measurement is viewed nonclassically, classical behavior can be maintained only if measurements of the system's behavior are made at low resolution. Explosive growth in the number of states that a system can be interpreted as traversing is thus prevented. No physical system can, therefore, be interpreted as computing every function.

The nonclassical behavior of a dynamical system can be described by a function that specifies the transition probabilities between every pair of states. Such a function can be calculated straightforwardly by a von Neumann machine; hence continuous dynamical systems do not have greater computational power than a universal Turing machine even in the nonclassical domain.

These results illustrate the general importance of physical considerations in the algorithmic description of physical systems. While such considerations are clearly relevant to engineering and implementation issues, they are also relevant to the theory of computation as a behavior of physical systems, and to the computational interpretation of such systems. Physical considerations are universal: it is significant that nowhere in the above discussion was any assumption made that the behaviors in question intuitively qualify as "computations." The neglect of universal physical principles, such as the nonclassical nature of measurement, can lead to the generation of artifactual paradoxes such as that of "underlying dynamics," and can obscure important features of computation as a physical process.

Acknowledgements.

I thank Srinivas Kankanahalli and Carol Soderlund for many discussions of the issues raised in this paper. This work was partially supported by NASA Innovative Research Program grant NAGW-1592 to the author and J. Barnden.

Notes.

1. "State space" is the standard term for the space of all possible machine configurations in AI (e.g. Barr and Feigenbaum, 1981, Ch. IIB). This usage of "state space" denotes the configuration space or ID space of a machine, as these terms are standardly used in computer science (e.g. Hopcroft and Ullman, 1979, Ch. 7), not the space of states of the control.
2. Many, from Einstein onward, have argued that the physical world may, after physics has progressed beyond its current stage, turn out to be classical afterall. This argument is typically based on aesthetic considerations, and all testable predictions of alternative, classical theories in physics have thus far proven false (see the references to Greenberger noted in the text). Our best description of the physical world has become progressively less classical since the development of field theory in the 19th century, and there is no reason, other than wishful thinking, to believe that this trend will not continue as experimental techniques improve.

References.

- Ashby, W. R. (1956) *An Introduction to Cybernetics* (London: Chapman and Hall).
- Barr, A. & Feigenbaum, E. (1981) *The Handbook of Artificial Intelligence*, Vol. 1 (Los Altos, CA: Kaufmann).
- Cohen, M. S. (1984) Interacting nonlinear waves in a neural continuum model: Associative memory and pattern recognition. In C. Rogers and T. B. Moodie (eds.) *Wave Phenomena: Modern Theory and Applications* (Amsterdam: Elsevier), 337-347.
- Fields, C. & Dietrich, E. (1987) A stochastic computing architecture for multi-domain problem solving. In Z. Ras, M. Emrich, and A. Jankowski (eds.) *Proceedings of the Second International Symposium on Methodologies for Intelligent Systems*, Colloquim Volume (Oak Ridge National Laboratory: ORNL-6417), 227-238.
- Freeman, W. J. (1988) Dynamic systems and the "subsymbolic level," *Behavioral and Brain Sciences*, 11: 33-34.
- Greenberger, D. M. (ed.) (1986) *New Techniques and Ideas in Quantum Measurement Theory* (New York: New York Academy of Sciences).
- Grossberg, S. (1988) Nonlinear neural networks: Principles, mechanisms, and architectures, *Neural Networks*, 1: 17-61.
- Hopcroft, J & Ullman, J. (1979) *Introduction to Automata Theory, Languages, and Computation* Reading, MA: Addison-Wesley).
- Jauch, J. M. (1968) *Foundations of Quantum Mechanics* (Reading, MA: Addison-Wesley).
- Kohonen, T. (1988) *Self-organization and Associative Memory*, 2nd Ed. (Berlin: Springer).
- Smolensky, P. (1988) On the proper treatment of connectionism, *Behavioral and Brain Sciences*, 11: 1-23.