© 2004 Mind and Matter Vol. 2(2), pp. 29–51

# Incompatible Implementations of Physical Symbol Systems

*Peter beim Graben*
*Institute of Linguistics and*
*Interdisciplinary Center for the Dynamics of Complex Systems*
*University of Potsdam, Germany*

**Abstract**

Classical cognitive science assumes that intelligently behaving systems must be symbol processors that are implemented in physical systems such as brains or digital computers. By contrast, connectionists suppose that symbol manipulating systems could be approximations of neural networks dynamics. Both classicists and connectionists argue that symbolic computation and subsymbolic dynamics are incompatible, though on different grounds. While classicists say that connectionist architectures and symbol processors are either incompatible or the former are mere implementations of the latter, connectionists reply that neural networks might be incompatible with symbol processors because the latter cannot be implementations of the former. In this contribution, the notions of "incompatibility" and "implementation" will be criticized to show that they must be revised in the context of the dynamical system approach to cognitive science. Examples for implementations of symbol processors that are incompatible with respect to contextual topologies will be discussed.

## 1. Introduction

Classical cognitive science relies upon the *Physical Symbol System* (PSS) hypothesis (Newell and Simon 1976) that intelligent behavior is essentially symbol processing which is in some way *implemented* by the dynamics of a physical system. Cognitive neuroscience is actually concerned with the question how these PSS's are implemented in the functioning of neural networks in the brain (Gazzaniga 2000). These attitudes are common among researchers in the field. However, there is an ongoing controversy whether neural network models could be regarded as alternatives to the symbolic paradigm providing proper *cognitive architectures* which are more than "mere" implementations (Smolensky 1991, p. 203) of symbol processors (Smolensky 1988, Fodor and Pylyshyn 1988, Smolensky 1991, Fodor and McLaughlin 1990, Chalmers 1990, Fodor 1997, Prince and Smolensky 1997). With his *Proper Treatment of Connectionism* (PTC), Smolensky (1988, p. 2) claimed that connectionist neural networks are even more than physical implementations of symbol systems for they operate neither at the "conceptual level" as these do, nor at the "neural level" of physiology (Smolensky 1988, pp. 5f). They rather constitute a "subconceptual level" (Smolensky 1988, p. 6) where "subsymbols", namely activation vectors, evolve according to nonlinear differential equations (Smolensky 1988, pp. 6f). Higher "conceptual levels" emerge by interpreting distributed activation patterns across the units of the network as symbols and their interactions as symbol processing for which a "complete, formal, and precise description" (Smolensky 1988, p. 6) in terms of algorithmic computation is not feasible (Smolensky 1991, p. 203). Thus, Smolensky argued, the PSS and the PTC account were *incompatible* and the PSS hypothesis would even be an approximation of the proper connectionist dynamics (Smolensky 1988, p. 7).

This position was attacked by the classicists Fodor and Pylyshyn (1988), Fodor and McLaughlin (1990) and Fodor (1997) advocating the PSS hypothesis. They argued that cognitive processes have a complex "combinatorial syntax and semantics" where "atomic" symbols can be merged into "molecular" symbols and molecular symbols further into more complex ones (Fodor and Pylyshyn 1988, p. 12) thus yielding a hierarchy of symbol structures whose "constituents" have "causal roles" in mental reasoning (Fodor and McLaughlin 1990, p. 199). According to Fodor and Pylyshyn (1988, p. 10) connectionist architectures of the mind do not have constituents of "representational states" and are, hence, either incompatible with PSS's or they are mere implementations at the level of "*non*representational states" (Fodor and Pylyshyn 1988, pp. 24, 28, 39, 41, 50).

Against Fodor and Pylyshyn (1988), Fodor and McLaughlin (1990) and Fodor (1997), Smolensky (1991) and Chalmers (1990) objected that connectionist neural networks allow indeed for constituent structures in distributed representations. Smolensky (1991, pp. 212ff) proposed his "tensor product representations" which map a hierarchical symbol structure onto an activation vector spanned by tensor products of "filler" vectors and "role" vectors. This mapping is unique if fillers and roles form a basis of their own subspaces, a point that was not appropriately acknowledged by Fodor and McLaughlin (1990) and Fodor (1997) in their replies. However, they are right as regards the lack of causal efficacy of constituents in the tensor product representation, thus concluding again that connectionist models are incompatible with symbolic ones as *architectures of the mind*.

More recently, Prince and Smolensky (1997) have shown that the symbolic architecture of *Optimality Theory* (OT) is compatible with connectionist *harmony machines* (Smolensky 1986) because OT is a "qualitatively different formal system at a higher level of analysis" of the subsymbolic dynamics of a neural network (Prince and Smolensky 1997, p. 1606).

Then, OT must be regarded as being implemented by the network's dynamics according to the definition of Smolensky (1988, pp. 59f). Therefore, we know examples of compatible implementations of PSS's; another one is discussed by Blutner (2004).

Today, connectionism is largely assimilated by the *Dynamical Systems Approach* (DSA) to cognition (Kelso 1995, van Gelder 1998, Beer 2000, beim Graben *et al.* 2004b) which is very successful in describing non-symbolic processes such as motor control or phonetics (Kelso 1995, Jirsa *et al.* 1998, Raçzasek *et al.* 1999). In this new context, questions arise as to the sense in which DSA models are incompatible with PSS's and whether they are implementations of the latter.

The aim of this paper is to discuss these issues. In order to do so, I shall criticize the concepts *incompatibility* in Sect. 2 and *implementation* in Sect. 3 with respect to dynamical systems and symbol processors. It will be shown in Sect. 2 that the concept of incompatibility was not appropriately used by both Fodor and Pylyshyn (1988) and Smolensky (1988) and must therefore be revised. In Sect. 3, I will then provide three examples for DSA models which are in fact implementations of PSS's in the sense of Smolensky (1988, 1991), but which are incompatible with other PSS's and with the underlying dynamics.

## 2. Incompatibility

The concept of incompatibility derives originally from quantum physics (Raggio and Rieckers 1983) where two observables are called *incompatible* if they are not precisely measurable simultaneously. Any precise measurement of one of them inevitably prevents a precise measurement of the other. This can be formally expressed by the notion of an *eigenstate*: A physical state is an eigenstate of an observable if the observable is *dispersion-free* with respect to the given state, i.e. if a measurement of the observable always yields the same precise measurement result. Two observables are *compatible* if all eigenstates of one of them are also eigenstates of the other one (and *vice versa*) and if their eigenstates span the whole state space (beim Graben and Atmanspacher 2004). Observables which are not compatible are called incompatible. In this case two observables do not share all eigenstates, and if one observable is dispersion-free in its eigenstate then the other one is not dispersion-free in the same state, and thereby not precisely measurable. Observables are maximally incompatible if they have no common eigenstate. Such observables are called *complementary*, which is the typical situation in quantum physics (beim Graben and Atmanspacher 2004).

The notion of a precise description also plays an important role in Smolensky's definition of the "incompatibility of the symbolic and the subsymbolic paradigms" (Smolensky 1988, p. 7) when a "subconceptual, connectionist dynamical system [...] does not admit a complete, formal, and precise conceptual-level description" (Smolensky 1988, p. 7), since such a description is feasible at the subconceptual level only (Smolensky 1988, pp. 6f). What is a "complete, formal, and precise description" in this context? By settling this question, I shall show how the concept of incompatibility should be revised for a DSA to cognition.

Let us first turn to the notion of a "formal" description. As Smolensky (1988, p. 6) explained, "there will generally be no precisely valid, complete, computable formal principles at the conceptual level; such principles exist only at the level of individual units — the *subconceptual level*." And Smolensky (1991, p. 203) added that "mental representations and mental processes are *not* supported by the same formal entities – there are no 'symbols' that can do both jobs. The new cognitive architecture is fundamentally two-level: formal, algorithmic specification of processing mechanisms, on the one hand, and semantic interpretation, on the other, must be done at two different levels of description."

These quotations illustrate what Smolensky had in mind: a "formal description" consists of "computable formal principles" allowing for the "formal, algorithmic specification of processing mechanisms". That is, a *formal description* should be a *mathematical model of a dynamical system* in its widest sense, including algorithmic or stochastic descriptions. We shall start with *deterministic dynamical systems* and take stochasticity into account later.

## 2.1 Dynamical Systems

Deterministic dynamical systems are described by a time-dependent *state* $\mathbf{x}(t)$ evolving according to deterministic differential or difference equations along a *trajectory* exploring the *phase space* $X$, i.e. the set of all possible states. The set of possible trajectories as solutions of the dynamical equations defines the *flow map* $\Phi^t : X \to X$ of the system such that for any given *initial condition* $\mathbf{x}(0)$ at time $t = 0$, $\mathbf{x}(t) = \Phi^t(\mathbf{x}(0))$ is its successor at a later time $t$.

A neural network consists of $n$ *units* which might be more or less activated, thus assuming real values between, say, 0 and 1, where 0 means no activation and 1 maximal activation. The state of the network is therefore given by an activation vector of dimension $n$, such that the phase space is the hypercube $X = [0, 1]^n \subset \mathbb{R}^n$. Units change their activation according to the connectivity of the network. A unit receives either excitatory or inhibitory input from other units, thereby changing its activation by the influence of the weighted sum of all other units that are connected to it. This gives rise to the dynamics which is usually described by nonlinear differential or difference equations.

Following Smolensky (1988), the DSA to a connectionist model provides the "subconceptual level" of a cognitive architecture. The deterministic differential equations admit a "formal" description of the "subsymbolic" dynamics (Smolensky 1988, pp. 6f). A formal description of a dynamical system is *complete* if all *relevant* variables are taken into account (Atmanspacher and Kronz 1999, p. 279). Here, relevance is decided by means of the formal model chosen to describe the system: relevant are the variables contained in the dynamical equations of the model. Consider again a neural network of $n$ units. The state of the system is an activation vector in $[0, 1]^n$ which evolves according to a differential equation with $n$ variables. Therefore, a complete description of the system refers to these $n$ relevant state variables. For instance, adding $m$ functions of the $n$ state variables to a description results in redundant and, hence, irrelevant variables. By contrast, a description with less than the $n$ relevant variables, is equivalent to a projection of the phase space onto a subspace such that the projected trajectories may intersect each other, thus violating determinism. Such a description is necessarily incomplete.

Considering that the formal mathematical model of a system determines its complete description, it is at least questionable whether such a description is possible for a physically implemented system. In this case one has to perform measurements on the system. In the most general mathematical framework which allows the treatment of dynamical systems, the so-called *algebraic quantum theory* (Haag 1992, Primas 1990), a measurement is described by a function $f : X \to \mathbb{R}$, a so-called *observable* mapping states in the phase space onto real numbers, the measurement results.[1] A complete description of a neural network of $n$ units can be obtained, e.g., by the simultaneous measurement of the $n$ observables $f_i(\mathbf{x}) = x_i$, $i = 1, \ldots, n$, which project the state $\mathbf{x}$ onto the $i$-th coordinate axis of the phase space. For particular physical systems such a complete set of observables can be obtained by virtue of the famous embedding theorem where a state vector is generated by delayed measurements of one experimental observable (Takens 1981).

Most *epistemic observables* (beim Graben and Atmanspacher 2004) assign the same measurement result to many different individual states in phase space, which therefore become *epistemically indistinguishable*. In other words, if $\mathbf{x}, \mathbf{y} \in X$ are two different states and $f : X \to \mathbb{R}$ is an observable such that $f(\mathbf{x}) = f(\mathbf{y})$ then we are unable to tell whether the system is in state $\mathbf{x}$ or in state $\mathbf{y}$ if we measure only $f$. Beim Graben and Atmanspacher (2004) call such states $\mathbf{x}, \mathbf{y}$ *epistemically equivalent with respect to the observable $f$*. Epistemic equivalence endows the phase space $X$ of a dynamical system with an equivalence relation entailing a

[1] Generally speaking, observables are complex-valued functions over the phase space. Physically meaningful observables are self-adjoined and, hence, real-valued functions.

*partitioning* of $X$ into pairwise disjoint subsets of epistemically equivalent states. We shall restrict ourselves to partitions of $X$ into a finite collection of subsets $\mathcal{P} = A_1, \dots, A_I$ $(I \in \mathbb{N})$.

Epistemic equivalence is not the only problem when dealing with complex physical systems. Physical measurements usually disturb the system to be measured. This uncontrollable impact of the environment upon a particular system must be treated by a stochastic approach: measuring the observable $f$ in state $\mathbf{x}$ yields realizations of a random variable $F$ which scatters around $f(\mathbf{x})$ (Atmanspacher and Primas 2003, p. 207). Then, even if there were no state $\mathbf{y}$ that is epistemically equivalent with $\mathbf{x}$, it would not be possible to know $\mathbf{x}$ since the pre-images of the realizations of $F$ under $f$ form a set of states scattering around $\mathbf{x}$. In practice, one computes the average over all outcomes of the disturbed observable, or, what is mathematically equivalent, one computes an ensemble average of all $f(\mathbf{y})$ over the whole phase space $X$ weighed by a particular probability density function (p.d.f.) $\rho$. These p.d.f.'s are called *statistical states* of the dynamical system in contrast to the *individual states* $\mathbf{x}$, which are points in the phase space.

## 2.2 Ontic and Epistemic Descriptions

Smolensky distinguished between the "conceptual level" where the dynamics of a connectionist dynamical system is interpreted in terms of symbol processing and the "subconceptual level" of the underlying "subsymbolic" evolution of activation patterns (Smolensy 1988, pp. 3ff). He pointed out that this distinction resembles that between the "microlevel" of quantum physics and the "macrolevel" of Newtonian mechanics (Smolensky 1988, pp. 12, 20) or even between the microlevel of statistical physics and the macrolevel of thermodynamics (Smolensky 1988, p. 11, see also Blutner 2004) insofar as macrophysical descriptions are only approximations of microphysical descriptions (Smolensky 1988, pp. 60f). However, this is not the case. Newtonian mechanics and thermodynamics are exact in their respective domains. We shall discuss in Sect. 2.3 in which sense the terms "approximation" and Smolensky's "precise description" should be applied.

Instead of talking about microlevels and macrolevels, Atmanspacher and Primas (2003) suggested to distinguish *ontic* and *epistemic* descriptions of physical systems. Epistemic descriptions refer to the "knowledge that can be obtained about an ontic state" (Atmanspacher and Primas 2003, p. 305, see also Atmanspacher 2000, p. 465), whereas ontic descriptions are complete with respect to a formal model. It turns out that ontic descriptions of a classical dynamical system refer to the individual points in phase space, whereas epistemic descriptions refer to statistical states as probability distributions over the phase space and to epistemic observables partitioning the phase space into sets of epistemically equivalent (ontic) states. Both aspects are closely connected because epistemic observables are defined by a certain statistical reference state providing a particular *context* (Primas 1990, beim Graben and Atmanspacher 2004).

Ontic and epistemic states do generally not coincide with the micro- and macrostates of statistical physics and thermodynamics (Atmanspacher 2000, p. 467). However, for classical dynamical systems, the individual, ontic states can be identified with the thermodynamic microstates. Let us consider the microcanonical ensemble of thermodynamics. Here the energy partitions the phase space of the system into (infinitely many) manifolds of epistemically equivalent states. The statistical state of the microcanonical ensemble for a given energy $E$ is then defined by the uniform p.d.f. over the surface assigned to the energy $E$. With respect to this distribution, all states of the manifold contribute the same statistical weight to the ensemble average of every observable. This means that any two states $\mathbf{x}, \mathbf{y}$ belonging to the microcanonical distribution are *statistically equivalent*. A macrostate of the microcanonical ensemble is therefore a set of energy-epistemically equivalent states which are also statistically equivalent.

Let $f$ be an epistemic observable that induces a finite partition $\mathcal{P}_f = A_1, \ldots, A_I$ $(I \in \mathbb{N})$ of the phase space $X$ into $I$ classes $A_i$ of states which are epistemically equivalent with respect to $f$. Then we can assign $I$ macrostates $\rho_i$ to the cells $A_i$ of the partition, which support the corresponding uniform p.d.f.'s. For any macrostate $\rho_i$ we can define epistemic observables $g_i(\mathbf{x}) = 1$ if $\mathbf{x} \in A_i$, and $g_i(\mathbf{x}) \neq 0$, $g_i(\mathbf{x}) \neq 1$ elsewhere. It then turns out that $\rho_i$ is an eigenstate of $g_i$. On the other hand, by construction $\rho_i$ is no eigenstate of any other $g_j$ $(j \neq i)$. Therefore two $g_i, g_j$ are incompatible. Are the two observables $g_i, g_j$ also maximally incompatible and hence complementary? The answer is: not necessarily, because any individual point in phase space defines an ontic state which is an eigenstate of any observable. However, complementary observables might appear in complex nonlinear dynamical systems where ontic states are *unphysical states* in the sense that they are not *epistemically accessible* (beim Graben and Atmanspacher 2004).

### 2.3 Symbolic Dynamics

So far we have discussed instantaneous descriptions of physical systems. Next we consider their dynamics. Let us restrict ourselves to time-discrete dynamical systems. Their flow is generated by a map $\Phi : X \to X$, such that a present state $\mathbf{x}(t) \in X$ at time $t$ is mapped onto its successor $\mathbf{x}(t+1) = \Phi(\mathbf{x}(t))$. A (discrete) trajectory is then obtained by iterating the map $\Phi$ recursively: $\Phi^t = \Phi^{t-1} \circ \Phi$. The map $\Phi$ may be invertible or not. While beim Graben and Atmanspacher (2004) discuss the special case of invertible flows, we shall here consider non-invertible flows since the examples presented in Sect. 3 belong to this class. Figure 1 shows an example for such a system where the point sequence $(\mathbf{x}(0), \mathbf{x}(1), \mathbf{x}(2), \mathbf{x}(3), \mathbf{x}(4))$ is a trajectory with initial condition $\mathbf{x}(0)$ ending in an asymptotically stable fixed point attractor $\mathbf{x}(4)$.

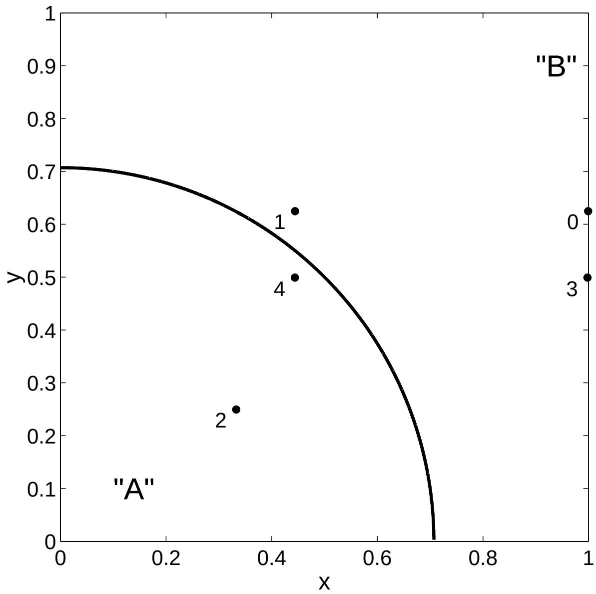


Figure 1: The origin of a symbolic dynamics: The phase space $X$ (the unit square) is partitioned into two disjoint subsets $A$ and $B$. A transient trajectory of a time-discrete dynamical system, the point sequence $(\mathbf{x}(0), \mathbf{x}(1), \mathbf{x}(2), \mathbf{x}(3), \mathbf{x}(4))$, ending in the fixed point $\mathbf{x}(4)$, is mapped onto a sequence of letters "BBABA".

Let $f$ be an observable of the system. We define a *continuous measurement of duration* $T$ by recording a *time series* of observations $\{f(t) = f(\Phi^t(\mathbf{x}(0))) \,|\, t = 0, \ldots, T-1\}$ starting with an initial condition $\mathbf{x}(0)$. From this record we may formally regain a point sequence in phase space $\{\mathbf{x}(t) = f^{-1}(f(t)) \,|\, t = 0, \ldots, T-1\}$ provided that the observable $f$ is one-to-one. A possible set of initial states $\Phi^{-t}(f^{-1}(f(t)))$ is then obtained by determining the pre-image under the flow in phase space. For invertible systems this set contains only a single state $\mathbf{x}(0)$. This is the proper meaning of a *precise description* of a dynamical system, where a complete description is required to enable a pointwise evaluation of states. Thus, a precise description refers to a complete description, which in turn depends on a formal description of a system: Smolensky's criteria for incompatibility are not independent. In the generic case of incomplete descriptions given by epistemic observables or non-invertible flows, the pre-images of the measured values of a time series under $f$ are non-singleton sets rather than individual points, $\{X(t) = f^{-1}(f(t)) \subset X \,|\, t = 0, \dots, T-1\}$, and their pre-images under the flow are also non-singleton sets, $\Phi^{-t}(f^{-1}(f(t))) \subset X$.

Returning to the epistemic observable $f$ inducing a finite partition $\mathcal{P}_f = A_1, \dots, A_I$ ($I \in \mathbb{N}$) of the phase space $X$, we introduce the basic concepts of *symbolic dynamics*. Measuring $f$ in the state $\mathbf{x}(0) \in X$ yields the value $f(0)$ of the time series, which is the same for all states $\mathbf{y}$ that are epistemically equivalent to $\mathbf{x}(0)$ with respect to $f$. That is, measuring $f$ in state $\mathbf{x}(0)$ tells that $\mathbf{x}(0)$ belongs to a particular class $A_{i_0}$. After one time step, the measurement of $f$ yields that the system is in state $\Phi(\mathbf{x}(0)) \in A_{i_1}$. Thus, we know that the initial state was in $A_{i_0} \cap \Phi^{-1}(A_{i_1})$. A continuous measurement of $f$ leads then to the estimation $\mathbf{x}(0) \in \bigcap_{t=0}^{T-1} \Phi^{-t}(A_{i_t})$. Given the partition $\mathcal{P}_f$, the sets $\bigcap_{t=0}^{T-1} \Phi^{-t}(A_{i_t})$ also partition the phase space for all combinations of the $A_{i_t}$ and for all measurement durations $T$. The resulting partitions are called *refinements* of $\mathcal{P}_f$.

By fixing the sequence of cells that are visited by a trajectory of the system, we obtain a one-sided sequence $s = a_{i_0} a_{i_1} a_{i_2} \dots$ of symbols $a_{i_t} \in \mathbf{A}_f$ where the index set of the partition is interpreted as a finite alphabet $\mathbf{A}_f$ of cardinality $I$. The initial state $\mathbf{x}(0)$ is thereby mapped onto a string $s$. Figure 1 shows how a binary partition of the unit square $X$ into the cells $A$ and $B$ leads to a sequence of symbols "BBABA" corresponding to the sets the trajectory $(\mathbf{x}(0), \mathbf{x}(1), \mathbf{x}(2), \mathbf{x}(3), \mathbf{x}(4))$ is exploring. Symbolic dynamics deals with the statistical and grammatical properties of such strings (Lind and Marcus 1995, Hao 1989). The mapping $\pi : \mathbf{x}(0) \mapsto s$ is generally not invertible. That is, the pre-image of a string $s$, $\pi^{-1}(s)$, is a non-singleton set $R$ rather than a point. These pre-images provide the finest partition of $X$ that is induced by the observable $f$ and its corresponding partition $\mathcal{P}_f$. The open sets of this *maximal refinement* constitute a *contextual topology* introduced by $\mathcal{P}_f$ (beim Graben and Atmanspacher 2004, Atmanspacher 2000).

Particular dynamical systems possess *generating partitions* whose refinements allow to approximate the individual points in the phase space by continuous measurements with arbitrary precision. Their contextual topologies are thus identical with the topology of the phase space (up to sets of vanishing probability measure). For generating partitions, the map $\pi^{-1}$ is invertible such that individual points in phase space are *epistemically accessible* by continuous measurements with increasing duration (beim Graben and Atmanspacher 2004). Moreover, the phase space dynamics and the symbolic dynamics are *topologically equivalent* for generating partitions. Beim Graben and Atmanspacher (2004) propose to call two different partitions *compatible* if they are both generating. If they are not generating, they are called *incompatible*. Compatible partitions allow for both compatible and incompatible observables, while incompatible partitions allow only for incompatible and complementary (i.e. maximally incompatible) observables since not every eigenstate of an observable is epistemically accessible by the dynamical refinement of the partition under continuous measurement (beim Graben and Atmanspacher 2004).

Non-generating partitions have topologies which are coarser grained than the topology of the phase space, so that they are incompatible with each other and incompatible with the structure of the phase space. This is presumably the property of incompatibility which Smolensky had in mind when he stated the incompatibility of the symbolic and the subsymbolic paradigms.

## 3. Implementations

The preceding section presented Smolensky's criteria for a "complete, formal, and precise description" of a dynamical system. We saw that a formal description defines the relevant variables from an ontic point of view, thus giving rise to a complete description. Correspondingly, a complete description is necessary for an ontic evaluation of observables, thereby defining a precise description. If there is an ontic and an epistemic, contextual account of a dynamical system, the epistemic description is neither complete nor precise with respect to the ontic description unless a generating partition is introduced by the chosen epistemic observables such that the two descriptions are compatible. In this case it is possible to approximate the ontic states with arbitrary precision by longer and longer lasting continuous measurements, thus leading to a complete epistemic description.

However, things become completely different if a new formal description is introduced also for the epistemic setup. Then one obtains a higher-level relative ontology (Atmanspacher and Kronz 1999, Atmanspacher and Primas 2003) where the formal model redefines the notions of completeness and precision. This is the case for the micro-macro relation between quantum mechanics and Newtonian mechanics or between statistical physics and thermodynamics. Regarding the microtheories, both macrotheories are approximations in the finer topologies. But their respective formal frames define the domains in which the macrotheories are complete and precise concerning the coarser contextual topologies. Consider the microcanonical ensemble of thermodynamics again. A macrostate is a uniform p.d.f. defined over an energy surface in the phase space. From the ontic perspective of the time evolution of individual points, a macrostate bears maximal uncertainty about a particular microstate. On the other hand, the evolution of p.d.f.'s considered as ontological objects is governed by deterministic equations, thus leading to another complete, formal description.

A similar micro-macro distinction can be applied to the conceptual and the subconceptual levels describing particular connectionist dynamical systems. Smolensky claimed that a complete, formal, and precise macrolevel description would be an *implementation* of a system with a complete, formal, and precise microlevel description (Smolensky 1988, p. 59). He called two descriptions of a connectionist system incompatible if one is complete, formal and precise at the subconceptual level, but the other one is only approximative at the conceptual level. The recently proposed approaches to optimality theory (Prince and Smolensky 1997) and non-monotonic logic (Blutner 2004) must be referred to as *compatible implementations* of PSS's by dynamical systems. By contrast, the next subsections will present three examples of nonlinear dynamical systems which are instances of implementations of PSS's and cannot be described by incompatible accounts in the sense of Fodor and Pylyshyn (1988) or Smolensky (1988). However, their epistemic descriptions will be incompatible regarding the coarse-grained topologies.

### 3.1 Example 1

Figure 2 shows five time series observed by a numerical experiment with a nonlinear dynamical system. The initial conditions were randomly chosen according to a uniform probability distribution defined over a region of the phase space which is actually the unit square depicted in Fig. 1.

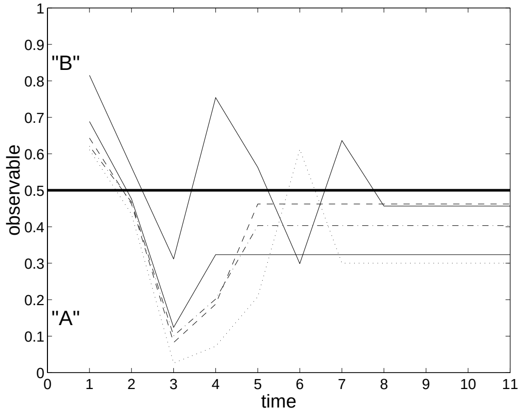


Figure 2: Five time series of an observable $f(\mathbf{x})$ for the same dynamical system as in Fig. 1, starting at randomly prepared initial conditions. The horizontal bold line denotes a partitioning of the range of the observable into the sets $A : f(\mathbf{x}) < 0.5$ and $B : f(\mathbf{x}) \geq 0.5$ which are the images of the cells $A$ and $B$ from Fig. 1 under the action of $f$.

One recognizes that all time series become eventually constant, thus indicating the existence of asymptotically stable fixed point attractors. The value zero in the range of the observable seems to correspond to an unstable, repelling fixed point since the dotted line diverges exponentially from the baseline. A symbolic dynamics can be introduced by any partition of the range of the observable (beim Graben *et al.* 2000). The theory of nonlinear time series analysis provides heuristics to find appropriate partitions (Wackerbauer *et al.* 1994, Daw *et al.* 2003). A first attempt might be a partition into cells of equal size. In this sense, the bold horizontal line in Fig. 2 denotes a threshold splitting the unit interval into the sets $A = [0, 0.5]$ and $B = ]0.5, 1]$. The pre-images of these sets under the particular observable are those regions in Fig. 1 that are labeled by the symbols $A$ and $B$. Figure 3 visualizes the symbolic dynamics of 50 time series (epochs) starting with randomly prepared initial conditions where black pixels denote the symbol “A” and white pixels denote the symbol “B”. The first five sequences are the encoded time series shown in Figure 2.

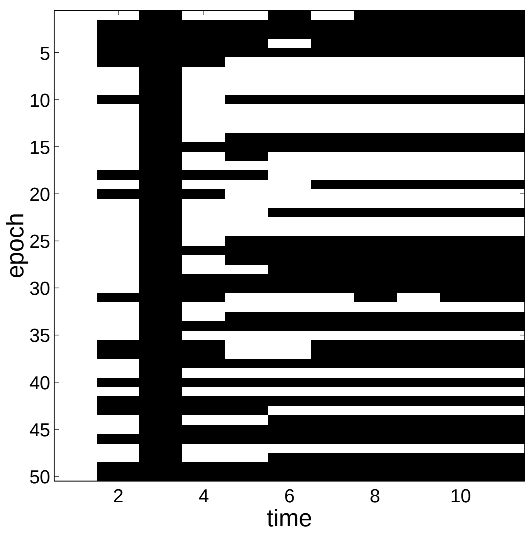


Figure 3: Symbolic dynamics of an ensemble of 50 time series (epochs) with random initial conditions of the same partitioned dynamical system as in Figs. 1 and 2. Here, black pixels denote the symbol “A” and white pixels “B”. The first five sequences are those of Figure 2.

It is easy to recognize that all initial conditions are situated in region $B$ of Figs. 1 and 2, respectively. After the first iteration all states are spread over the whole phase space but concentrated in region $A$ after the third iteration. The fourth iteration yields a random distribution of symbols again. Finally, after the fifth step, the states reach their attractors either in $A$ or in $B$. The particular context chosen by this partition is not very instructive yet. We shall consider this example as one of many possible *physical* contexts which might provide insight into the underlying dynamics.

Figure 4 unveils what is actually going on. The dynamics is defined by a piecewise affine linear map that comprises parallel translations, stretching and squeezing along the coordinate axes of the phase space. The domains of definition are the six rectangles, denoted $A$ to $F$ in Fig. 4. The flow acts in the following way: the map is simply the identity at the white rectangles $B$ and $C$, which are thereby mapped onto themselves, respectively. The light-gray rectangles $A$ and $D$ are expanded in both directions and shifted such that their bottom-left vertices coincide with the origin of the unit square. Their images are the whole unit square. The dark-gray rectangles $E$ and $F$ are shifted and squeezed along the $x$-axis, thereby being mapped onto the anthracite rectangles contained in $C$ and $D$, respectively.

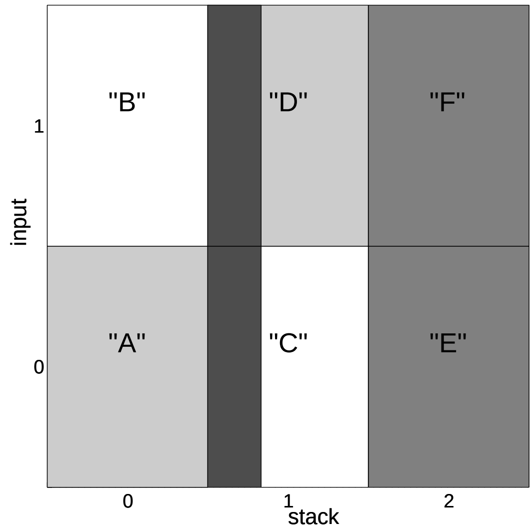


Figure 4: The *intended* partition of the dynamical system from Figs. 1–3 is given by the six rectangles $A$ to $F$. The dynamics is defined by piecewise affine linear maps at these domains.

From Fig. 4 we can deduce that the dynamics is asymptotically multistable. All states will be transiently rinsed into the rectangles $B$ or $C$ either, where each state is a fixed point. Therefore, all sequences of the six symbols $A$ to $F$ become eventually periodic either ending with $B$ or

$C$. It follows that the system does not have a generating partition (for a proof see the Appendix), hence all particular partitions are incompatible with each other. The symbolic dynamics is not topologically equivalent with the ontic dynamics of individual states.

The partition displayed in Fig. 4 is called the *intended partition* since it was explicitly used to define the piecewise affine linear dynamics. However, there is a further reason for this notion. Beim Graben *et al.* (2004a) have shown that the system presented in Figs. 1–5 has a straightforward formal interpretation as a PSS. The transients of the dynamics from time step 0 to 3 describe the state transitions of a pushdown automaton processing a sentence of a formal language (Hopcroft and Ullman 1979). This interpretation is not comprised at the ontic level of individual points hopping through the unit square. One has to consider rectangular macrostates moving around the phase space. Figure 5 shows the example underlying all previously discussed illustrations.

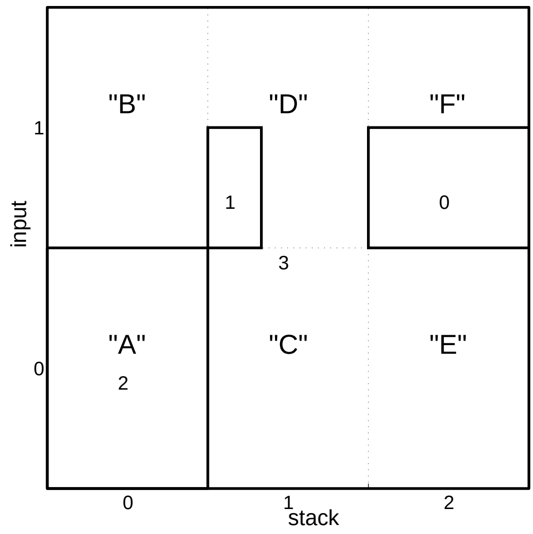


Figure 5: The transient evolution of an epistemic state of uniformly distributed initial conditions supported by the rectangle $R(0)$ through the unit square. The sequence of rectangles $(R(0), R(1), R(2), R(3))$ corresponds to the processing of a "sentence" by a pushdown automaton according to a particular context-free grammar. The last state $R(3)$ is the accepting state of the automaton.

We see that initial conditions which were randomly drawn from a uniform distribution supported by rectangle $R(0)$ are successively mapped onto states whose envelopes are the rectangles $(R(1), R(2), R(3))$. In the intended symbolic dynamics, this macroscopic trajectory is encoded by the sequence "FDA$\varepsilon$" where $\varepsilon$ denotes the final state $R(3)$ covering the whole unit square. Another representation of this sequence is more instructive. Projecting the cells of the partition onto the coordinate axes yields pairs of symbols "0", "1", and "2" which partition the $x$- and the $y$-axis. Then the sequence "FDA$\varepsilon$" becomes "$((2,1),(1,1),(0,0),(\varepsilon,\varepsilon))$" where we have replaced "$\varepsilon$" by the pair "$(\varepsilon,\varepsilon)$". This symbolic trajectory encodes both a constituental phrase structure tree and the parse of the processing automaton (beim Graben *et al.* 2004a). The formal language being processed is generated by a simple context-free grammar (Hopcroft and Ullman 1979) "$2 \rightarrow 10$" where "2" denotes the start symbol of the grammar whereas "1" can be interpreted as the SUBJECT and 0 as the PREDICATE of a sentence. The initial macrostate encodes the state 0 of the automaton having the start symbol "2" at the stack and the sentence "10", i.e. SUBJECT-PREDICATE, at the input tape. In the first step the parser "recognizes" the start symbol and "predicts" a SUBJECT according to the grammar by expanding the start symbol into the sequence "10". Now, the leading symbols at the stack and at the input tape are both "1" in state 1. The prediction was successful and the parser "attaches" the predicted SUBJECT to the SUBJECT found in the input by canceling both symbols from both tapes arriving in state 2. Here the parser finds again an agreement between the predicted PREDICATE at the stack and the PREDICATE provided by the input. A further attachment leads to the accepting state 3 where both tapes are empty.

This example supplies indeed an implementation, namely a complete, formal and precise conceptual and symbolic interpretation of a completely, formally and precisely described nonlinear dynamical system. However, the intended partition providing the symbol processing context is incompatible with any partition describing a physical context because none of the partitions is generating. Moreover, the symbolic and the subsymbolic descriptions are incompatible either, because there is no topological equivalence between them.

### 3.2 Example 2

My second example is adopted from Balkenius and Gärdenfors (1991) (see also Gärdenfors 1994, Blutner 2004). Balkenius and Gärdenfors (1991) discuss attractor neural networks (ANN) (Amit 1989) such as, e.g., Hopfield nets on an $n$-dimensional hypercube $X = [0,1]^n$ as their phase space, where a time-discrete flow $\Phi : X \rightarrow X$ mediates the dynamics. These connectionist dynamical systems are asymptotically (multi-)stable such that a *natural partition* of the phase space $X$ into the basins of attraction of the fixed point attractors exists. Again, these systems do not possess generating partitions by virtue of the proposition proven in the

Appendix. Hence, all particular partitions are incompatible with each other and with the topology of the phase space.

Balkenius and Gärdenfors (1991) introduce a conceptual level by particular hypercubic partitions, called *schemata*. A schema is given by a point $\mathbf{a} \in [0,1]^n$. The schema is said to be *accepted* by another point $\mathbf{x} \in [0,1]^n$ if $x_i \geq a_i$ for $i = 1, \ldots, n$ (Gärdenfors 1994). All accepted states $\mathbf{x}$ of a given schema $\mathbf{a}$ constitute a macrostate (a "cone") $A \subset X$ belonging to a partition $\mathcal{P}_{\mathbf{a}}$ that is induced by the schema $\mathbf{a}$. Among schemata a partial ordering is introduced by the relation $\mathbf{a} \succeq \mathbf{b}$ if $a_i \geq b_i$ for $i = 1, \ldots, n$, meaning that the schema $\mathbf{a}$ is *more specific* than $\mathbf{b}$. There is a minimal schema $\mathbf{0} = (0, \ldots 0) \in [0,1]^n$ and a maximal schema $\mathbf{1} = (1, \ldots 1) \in [0,1]^n$ such that the set of schemata forms a lattice. The respective lattice operations are given as $(\mathbf{a} \odot \mathbf{b})_i = \max(a_i, b_i)$ for the *conjunction* and $(\mathbf{a} \oplus \mathbf{b})_i = \min(a_i, b_i)$ for the *disjunction* of schemata $\mathbf{a}, \mathbf{b}$. Finally, a *complement* can be introduced by $(\mathbf{a}^*)_i = 1 - a_i$ establishing a De Morgan lattice as a provisional logic.

In order to define an inference relation, the dynamics has to be taken into account. Passing a schema $\mathbf{a}$ to the flow as an initial condition yields a trajectory of schemata $\mathbf{a}, \Phi(\mathbf{a}), \Phi^2(\mathbf{a}), \ldots, \hat{\mathbf{a}}$ ending at the attractor $\hat{\mathbf{a}}$. However, the fixed point $\hat{\mathbf{a}}$ is generally not more specific than the initial condition $\mathbf{a}$, in contrast to the intention of logical inference. Balkenius and Gärdenfors (1991) remedy this obstacle by *clamping* the time evolution of states. They introduce the clamped map $\Phi_{\mathbf{a}}(\mathbf{x}) = \Phi(\mathbf{x}) \odot \mathbf{a}$ and determine the fixed point $\hat{\mathbf{x}}_{\mathbf{a}} = \lim_{t\to\infty} \Phi^t_{\mathbf{a}}(\mathbf{x})$. Then, by $\hat{\mathbf{a}}_{\mathbf{a}} \succeq \mathbf{b}$ a non-monotonic inference relation $\mathbf{a} \rightsquigarrow \mathbf{b}$ is defined.

The latter example shows how the subsymbolic dynamics of an ANN can be interpreted as a non-monotonic logic at a formal conceptual level (Besnard *et al.* 2003). Again we have a formal, precise and complete account at each level and, thus, an implementation of a PSS. However, the symbolically interpreted epistemic states of the system belong to a lattice of incompatible partitions. There is no topological equivalence of the conceptual interpretation with the underlying dynamics.

### 3.3 Example 3

The concluding example brings us close to neurobiology. *Coupled Map Lattices* (CML) (Kaneko 1993) are becoming more and more popular modeling tools in computational neuroscience. They are spatially distributed lattices of $n$ vertices, where each vertex $k$ is occupied by a time-discrete dynamical system, e.g. the logistic map, at the local phase space $X_k = [0,1]$. The dynamics of the vertices is governed by their intrinsic flow $\Phi$ coupled to the states of particular other nodes (usually nearest neighbors). The evolution equation is then given as $x_k(t+1) = (1-\varepsilon)\Phi(x_k(t)) + (\varepsilon/m)\sum_{(j,k)} \Phi(x_j(t))$. The sum extends over those $m$

vertices $j$ which are connected to vertex $k$ (symbolically: $(j, k)$). The dynamics depends on two parameters, the coupling strength $\varepsilon$ and a control parameter $r$ determining the flow $\Phi$. For the logistic map, this dependence is given by $\Phi(x) = rx(1 - x)$.

Atmanspacher and Scheingraber (2004) investigated the stability properties of such CMLs in the fully chaotic regime of the individual vertex dynamics. The ordinary logistic map has a generating partition for $r = 4$ provided by the critical state $x_c = 0.5$: $A = [0, x_c], B = ]x_c, 1]$. Therefore, the phase space $X = [0, 1]^n$ of the CML of uncoupled logistic maps ($\varepsilon = 0$) is partitioned by $\mathbf{x}_c = (0.5, \ldots, 0.5)$ into hypercubes of equal size representing a generating partition as well. The symbolic dynamics of the CML can be regarded as a spatio-temporal pattern of "A"s and "B"s akin to that of Figure 3.

Increasing the coupling strength $\varepsilon$, Atmanspacher and Scheingraber (2004) have shown that a locally unstable fixed point $x_k^* = (r - 1)/r$ at each vertex becomes a global asymptotically stable fixed point for $\varepsilon \geq \varepsilon_c$, where the stabilization onset $\varepsilon_c$ depends on the network architecture and the updating procedure. If the coupling parameter is larger than this critical threshold, all lattice vertices arrive after a short transient phase at this attractor. Thus, the corresponding symbolic dynamics performs a simple algorithm which brings each vertex into the eventually periodic sequence of "B"s (or a white pixel in a visualization as in Fig. 3). Though the formal, algorithmic interpretation of that system as symbol processing is not very instructive, it shares an important property with the previous examples. After the global stabilization onset of the CML the system does not have a generating partition any more since it is asymptotically stable. We have a system which changes its coarse-grained contextual properties so drastically that any compatibility between partitions and the topological equivalence of any partition and the ontic dynamics is lost after the stabilizing bifurcation.

## 4. Discussion

The present study contributes to the long-lasting controversy between classicists defending the Physical Symbol System (PSS) hypothesis of cognitive science on the one hand and the connectionist/dynamicist camp on the other, arguing that symbolic computation is a higher-level interpretation of low-level nonlinear dynamics. Proponents of both lines of thought reasoned that symbol processing and dynamics were incompatible, though on different grounds. Classicists, such as Fodor and Pylyshyn, denied that dynamical systems could be proper cognitive architectures *of the mind* because these systems do not allow for combinatorial, constituental syntax and semantics with their commonly accepted properties of productivity, systematicity, compositionality, and causal efficacy of constituents.

However, they did not deny the possibility that symbolic processes are implemented by connectionist or dynamical systems which is in fact a necessary prerequisite for any *physical* symbol system.[2] The best known example is the implementation of a von Neumann machine by a digital computer. By contrast, connectionists such as Smolensky objected that the concept of an implementation was not properly used by the classicists, and that PSS's and connectionist dynamical systems are incompatible because the former are not mere implementations of the latter, but rather rough approximations.

I think that cognitive and computational neurosciences ought to seek implementations of PSS's by neurodynamical systems which are as formal, precise and complete as possible in their proper domains. I criticized Smolensky's earlier position that such descriptions are only feasible at the modeling instance. His definition of implementation entailed a category mistake because he did not observe that the conceptual and the subconceptual level of a system are treated by distinct formal accounts.

The three examples of implementations of PSS's that I have supplied bear a common resemblance. At the subconceptual level, the formal model introduces an ontological stance defining the concepts of completeness and precision. By partitioning the phase space of the ontic dynamics into epistemic macrostates, a coarse-grained contextual topology is provided through the maximal dynamically generated refinement of the partition. Regarding the ontic description, the epistemic account is neither complete nor precise and ontic states are only approximately accessible by epistemic means. On the other hand, as soon as a formal description is introduced at the conceptual level, this account determines a new relative ontology, thereby redefining complete and precise descriptions in this frame of reference. Symbolic dynamics supplies the necessary tools for such a formal treatment of epistemic states.

In the first example of Sect. 3.1, the intended partition of the unit square enables the interpretation of the transient dynamics of symbolic states as syntactic language processing. This is formally described by the theory of pushdown automata (Hopcroft and Ullman 1979), thus leading to an implementation of a PSS (in the sense of Smolensky) by a nonlinear dynamics. Interestingly, due to the construction of the system's flow, constituents, namely the leading symbols at stack and input tape *are* indeed causally efficacious. Therefore, the system is actually more than a mere implementation in the sense of Fodor and Pylyshyn. This holds for a large class of such systems since Moore (1990) has proven that any Turing machine can be represented by a piecewise affine linear map acting at the unit square. The second example presented in Sect. 3.2 reviews the interpretation of particular macrostates of an attractor neural network

[2] This point was not correctly taken into account by Chalmers (1990).

(ANN) in terms of non-monotonic logical reasoning. Again, a complete and precise formal description at the conceptual level and therefore an implementation according to Smolensky's criteria is achieved. The last example in Sect. 3.3 has no straightforward conceptual interpretation in terms of cognitive models. It simply performs a painter's work creating a homogeneous stable spatial pattern. Its high-level dynamics allows for a precise, complete and formal description, thereby providing an implementation of a simple PSS.

Furthermore, these examples illustrate that the concept of incompatibility has not been used appropriately by both classicists and connectionists. Two descriptions of a system are not simply incompatible if they are not implementations of each other. By contrast, the concept of incompatibility deeply refers to a topology that is connected to the chosen description. Originally introduced into quantum mechanics, where observables are incompatible if they are not precisely measurable simultaneously, beim Graben and Atmanspacher (2004) have generalized this concept to epistemic observables and partitions even of classical dynamical systems. They call two partitions compatible if they are arbitrarily well refinable by continuous measurements, eventually reaching the topology of the underlying ontic description. In this case, when the partitions are generating, an epistemic account yields approximately complete and precise descriptions of the system. Non-generating partitions are incompatible. They generate different topologies by the process of dynamic refinement which are coarser grained than the ontic topology. Systems with incompatible partitions cannot be consistently described by a distinguished epistemic account. Each conceptual and therefore epistemic description is principally different from any other. There is no *lingua franca*, no uniquely determined common formal account feasible for such systems.

This is demonstrated by the three examples discussed. They have in common that their dynamics is asymptotically stable, thereby preventing generating partitions. Example 3.1 allows for many physical contexts in terms of nonlinear data analysis corresponding to incompatible partitions. Among them is the intended partition upon which the construction of the dynamics relies. Only this particular partition allows for the interpretation of the system's behavior as language processing. The second example (Sect. 3.2) introduces symbolic schemata as partitions of the phase space of an ANN. These schemata, constituting a De Morgan lattice as a provisional logic, are all incompatible with each other. Moreover, no schema enables the approximation of individual activation patterns accepting it because the dynamics is not generating and therefore not topologically equivalent with the structure of the phase space. The last example (Sect. 3.3) illustrates that bifurcations in a nonlinear dynamical system may destroy a unique formal description by creating a host of incompatible epistemic descriptions.

Finally we can conclude that Smolensky's hypothesis (10) (Smolensky 1988, p. 7) has not to be rejected but rather to be adopted: "Valid connectionist models are merely implementations, for a certain kind of parallel hardware, of symbolic programs that provide exact and complete accounts of behavior at the conceptual level." Optimality Theory (OT; Prince and Smolensky 1997) and the reconstruction of the non-monotonic logic calculus of Balkenius and Gärdenfors (1991) by Blutner (2004) actually provide examples for such implementations.[3] However, these symbol systems are implemented by a symbolic interpretation of activation vectors in neural networks. They are therefore described by the ontological stance of individual points in phase space, which always allows for compatible descriptions. Thus, OT and Blutner's calculus must be regarded as compatible implementations of PSS's in the light of the appropriate definition of incompatibility. On the other hand, epistemic conceptual descriptions are incompatible with the subsymbolic dynamics if its corresponding partitions are not generating, thus preventing topological equivalence between the conceptual and the subconceptual level accounts. Moreover, different conceptual level descriptions will be incompatible with each other, hence preventing the existence of one uniquely determined distinguished formal account.

## Appendix

*Proposition.* An asymptotically (multi-)stable dynamical system does not possess generating partitions.

*Proof.* Consider a time discrete non-invertible dynamical system with phase space $X$ and flow $\Phi$. This asymptotically (multi-)stable dynamical system is naturally partitioned into the basins of attraction of its fixed point attractors. Call this partition $\mathcal{P}_{\mathrm{nat}} = \{B_1, B_2, \dots\}$ (it might be either finite or infinite). Let $\mathbf{x}^* \in X$ be an asymptotically stable fixed point with basin of attraction $B_k \in \mathcal{P}_{\mathrm{nat}}$. Let $\mathcal{P} = \{A_1, \dots A_I\}$ be an arbitrary finite partition of $X$ such that $\mathbf{x}^* \in A_1$ (possibly after a permutation of the alphabet). Then, $\mathbf{x}^*$ is mapped by the symbolic encoding $\pi$ onto the periodic sequence $1^\infty$. Choose now a neighborhood $R_{\mathbf{x}^*}$ of $\mathbf{x}^*$ such that $R_{\mathbf{x}^*} \subset B_k$. The whole set $R_{\mathbf{x}^*}$ is thereby mapped by $\pi$ onto the sequence $1^\infty$. Hence $R_{\mathbf{x}^*} \subset \pi^{-1}(1^\infty)$. That is $\pi^{-1}(1^\infty)$ contains more than exactly one element and $\pi$ is therefore not invertible. Thus, $\mathcal{P}$ is not generating.

[3]Note that OT itself has been described as a family of non-monotonic logics by Besnard *et al.* (2003.)

## Acknowledgements

I thank Carsten Allefeld and Harald Atmanspacher for stimulating discussions and the Deutsche Forschungsgemeinschaft for financial support within the research group on "Conflicting Rules in Cognitive Systems".

## References

Amit D. J. (1989): *Modeling Brain Function. The World of Attractor Neural Networks*, Cambridge University Press, Cambridge.

Atmanspacher H. (2000): Ontic and epistemic descriptions of chaotic systems. In *Computing Anticipatory Systems: CASYS 99*, ed. by D. M. Dubois, Springer, Berlin, pp. 465–478.

Atmanspacher H. and Kronz F. (1999): Relative onticity. In *On Quanta, Mind and Matter. Hans Primas in Context*, ed. by H. Atmanspacher, A. Amann and U. Müller-Herold, Kluwer, Dordrecht, pp. 273–294.

Atmanspacher H. and Primas H. (2003): Epistemic and ontic quantum realities. In *Time, Quantum, and Information*, ed. by L. Castell and O. Ischebeck, Springer, Berlin, pp. 301–321.

Atmanspacher H. and Scheingraber H. (2004): Inherent global stabilization of unstable local behavior in coupled map lattices. *International Journal of Bifurcation and Chaos*. In press.

Balkenius C. and Gärdenfors P. (1991): Nonmonotonic inferences in neural networks. In *Principles of Knowledge Representation and Reasoning*, ed. by J. A. Allan, R. Fikes, and E. Sandewall, Morgan Kaufmann, San Mateo, pp. 32–39.

Beer R.D. (2000): Dynamical approaches to cognitive science. *Trends in Cognitve Sciences* **4**, 91–99.

beim Graben P. and Atmanspacher H. (2004): Complementarity in classical dynamical systems. arXiv: nlin.CD/0407046.

beim Graben P., Jurish B., Saddy D., and Frisch S. (2004a): Language processing by dynamical systems. *International Journal of Bifurcation and Chaos* **14**, 599 –621.

beim Graben P., Saddy D., and J. Kurths, eds. (2004b): *Cognition and Complex Brain Dynamics*, *International Journal of Bifurcation and Chaos* **14** (special issue).

beim Graben P., Saddy J.D., Schlesewsky S., and Kurths J. (2000): Symbolic dynamics of event–related brain potentials. *Physical Review E* **62**, 5518–5541.

Besnard P., Fanselow G., and Schaub T. (2003): Optimality theory as a familiy of cumulative logics. *Journal of Logic, Language and Information* **12**, 153–182.

Blutner R. (2004): Nonmonotonic inferences and neural networks. *Synthese* **141**, in press.

Chalmers D.-J. (1990): Why Fodor and Pylyshyn were wrong: the simplest refutation. In *Proc. 12th Ann. Conf. Cognitive Science Society*, Lawrence Erlbaum, Hillsdale, pp. 340–347.

Daw C.S., Finney C.E.A., and Tracy E.R. (2003): A review of symbolic analysis of experimental data. *Review of Scientific Instruments* **74**, 915–930.

Fodor J. (1997): Connectionism and the problem of systematicity (continued): why Smolensky's solution still doesn't work. *Cognition* **62**, 109–119.

Fodor J. and McLaughlin B. P. (1990): Connectionism and the problem of systematicity: why Smolensky's solution doesn't work. *Cognition* **35**, 183–204.

Fodor J. and Pylyshyn Z. W. (1988): Connectionism and cognitive architecture: A critical analysis. *Cognition* **28**, 3–71.

Gärdenfors P. (1994): How logic emerges from the dynamics of information. In *Foundations of Computation*, ed. by J. v. Eijck and A. Visser, MIT Press, Cambrigde MA, pp. 49–77.

Gazzaniga M.S., ed. (2000): *The New Cognitive Neuroscience*, MIT Press, Cambridge MA.

Haag R. (1992): *Local Quantum Physics: Fields, Particles, Algebras*, Springer, Berlin.

Hao B.-L. (1989): *Elementary Symbolic Dynamics and Chaos in Dissipative Systems*, World Scientific, Singapore.

Hopcroft J.E. and Ullman J.D. (1979): *Introduction to Automata Theory, Languages, and Computation*, Addison Wesley, Menlo Park.

Jirsa V.K., Fuchs A., and Kelso J.A.S. (1998): Connecting cortical and behavioral dynamics: Bimanual coordination. *Neural Computation* **10**, 2019–2045.

Kaneko K., ed. (1993): *Theory and Applications of Coupled Map Lattices*, Wiley, New York.

Kelso J.A.S. (1995): *Dynamic Patterns. The Self–Organization of Brain and Behavior*, MIT Press, Cambrigde MA.

Lind D. and Marcus B. (1995): *An Introduction to Symbolic Dynamics and Coding*, Cambridge University Press, Cambridge. Reprint 1999.

Moore C. (1990): Unpredictability and undecidability in dynamical systems. *Physical Review Letters* **64**, 2354–2357.

Newell A. and Simon H. A. (1976): Computer science as empirical inquiry: Symbols and search. *Communications of the Association for Computing Machinery* **19**, 113–126.

Primas H. (1990): Mathematical and philosophical questions in the theory of open and macroscopic quantum systems. In *Sixty-two Years of Uncertainty*, ed. by A.I. Miller, Plenum Press, New York, pp. 233–257.

Prince A. and Smolensky P. (1997): Optimality: From neural networks to universal grammar. *Science* **275**, 1604–1610.

Raçzasek J., Tuller B., Shapiro L.P., Case P., and Kelso S. (1999): Categorization of ambiguous sentences as a function of a changing prosodic parameter: A dynamical approach. *Journal of Psycholinguistic Research* **28**, 367–393.

Raggio G.A. and Rieckers A. (1983): Coherence and incompatibility in $W^*$-algebraic quantum theory. *International Journal of Theoretical Physics* **22**, 267–291.

Smolensky P. (1986): Information processing in dynamical systems: Foundations of harmony theory. In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, ed. by D.E. Rumelhart, J.L. McClelland, and the PDP Research Group, MIT Press, Cambridge MA, pp. 194–281.

Smolensky P. (1988): On the proper treatment of connectionism. *Behavioral and Brain Sciences* **11**, 1–74.

Smolensky P. (1991): Connectionism, constituency, and the language of thought. In *Meaning in Mind. Fodor and his Critics*, ed. by B. Loewer and G. Rey, Blackwell, Oxford, pp. 201–227.

Takens F. (1981): Detecting strange attractors in turbulence. In *Dynamical Systems and Turbulence, Lecture Notes in Mathematics* **898**, ed. by D.A. Rand and L.-S. Young, Springer, Berlin, pp. 366–381.

van Gelder T. (1998): The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences* **21**, 615–628.

Wackerbauer R., Witt A., Atmanspacher H., Kurths J., and Scheingraber H. (1994): A comparative classification of complexity measures. *Chaos, Solitons & Fractals* **4**, 133–173.

*Received: 17 August 2004*
*Revised: 28 September 2004*
*Accepted: 4 October 2004*

*Reviewed by Reinhard Blutner and Gerda Demmel.*