

Behavioral States: Linking Functional and Physical Descriptions

Matthias Scheutz

Department of Computer Science and Engineering
University of Notre Dame, Notre Dame, IN 46556, USA
mscheutz@cse.nd.edu

Abstract

In this paper, I will introduce the notion of “behavioral state” as a means to bridge the gap between functional specifications and their implementations, borrowing from ethology and to some extent from research in behavior-based robotics. First, I will briefly sketch some of the problems resulting from mere functional descriptions for the designer of a mind. Then I will define the notion of behavioral state and locate its place as mediator between functional and physical states. After sketching the way in which I foresee behavioral states could be used in the design of minds, I will conclude that an intermediary level of architectural specification between functional and physical description will be of great advantage (if not necessity) in designing a mind, regardless of the success of the notion of behavioral state.

1 Introduction

Ever since Descartes, philosophers interested in the mind have divided the world into the mental and the physical realm contemplating exactly how these two realms are related. While this issue is still far from being resolved, today’s most commonly held view on the “mind-body” problem in the philosophy of mind is *functionalism*, i.e., the claim that mental states *are* functional states, which somehow “supervene” on physical states.¹ The general understanding is that mental states (i.e., states such as “believing that p ” or “desiring x ”, or even psychological predicates such as “pain” or “pleasure”) can be explained in terms of functional states and functional architectures.

Besides the fact that to my knowledge no one has attempted to account for concepts from folk psychology by specifying *in detail* a functional architecture for them, it seems to me that there will still be major obstacles for the artificial intelligence researcher who wants to build actual agents (that realize a given functional architecture). Even if such an architecture could be provided the question remains how functional states are related to physical states? Furthermore, should a

combination of “computational-physical” states be used as realizers instead of physical states alone? What constraints does the architecture impose on the implementing system?

Relating functional states *directly* to physical states is very unlikely to succeed in the light of multiple realization arguments for functional architectures (the more complex the architecture gets, the less we will be able to see what kinds of possibly very diverse physical systems will share the functional specification). The level of functional specification of the psychology of minds will be too high and abstract a level of description to suggest *possible implementations* of the functional states (not to mention all the problems connected with the involved notion of “implementation” or “realization” that seem to be largely ignored by the philosophical community).²

It is my conviction that functional specifications of psychologies are not sufficient to suggest ways to build a mind. To be of any practical importance in designing a mind at all, a level of description of a cognitive architecture has to incorporate at least *some* of the relevant physical properties of its possible implementations, which will constrain both possible implementations as well as functional architectures. In this article, I will suggest such an intermediary level, which I call *the level of behavioral states*. This level of description is largely inspired by ethological studies of animal behavior and to some extent by research in behavior-

¹ The questions of exactly how these states supervene on the physical and in what kinds of structures they are realized are rarely addressed in detail, let alone answered satisfactorily. This is most likely due to the fact that the notions of “realization” and “supervenience” are mostly used as unexplained “primitive” terms in the philosophical literature (which is quite surprising given the theoretical importance and practical consequences that hinge upon them). Although some have attempted more or less precise definitions of “realization”—e.g., Kim, Block, et al.—these definitions are not very helpful for those who, interested in building minds, are trying to understand the relation between architectures and their implementations.

² Note that this obviously does not hold for all functional specifications: a functional specification of an abstract finite state automaton, for example, can be easily related to physical states in a standard PC by “implementing” the automaton in a programming language.

based robotics and will therefore bear the signia of these intellectual sources very visibly on its sleeves.

First, I will briefly point to problems resulting from mere functional descriptions for the designer of a mind. Then I will introduce the notion of “behavioral state” and locate its place as mediator between functional and physical states. I will sketch the way I forsee that behavioral states could be used in the design of minds on a simple cognitive architecture. Finally, I shall conclude that regardless of the success of the notion of behavioral state in designing minds, an intermediary level of architectural specification between functional and physical description will be of advantage (if not necessity) in designing a mind.

2 Functionalism

2.1 The Functionalist Picture

A functional specification of a cognitive architecture consists a set of input states, a set of output states, and a set of “inner” or “functional” states together with a specification of how they are causally related. That way it is possible to determine what state a cognitive system will be in next, given the current state and all the input conditions.³ While input and output conditions have to be tied to physical inputs and outputs, the functional states do not require a direct correspondence to their physical realizers as expressed in the phrase that “functional states supervene on physical states” (e.g., see Kim, 1997). This lack of a “direct” correspondence between functional and physical states is what gives functionalism its explanatory power, while keeping it metaphysically palatable: it combines advantages of behavioristic approaches to mind (i.e., considering solely the input-output behavior of an organism) with advantages of identity theories (i.e., mental state/event tokens are physical state/event tokens) leaving out the pitfalls of both such as the lack of being able to account for “inner states” in the former, and the requirement of type identities between mental and physical state/event types of the latter. Yet, this strength comes at a price: it is not clear what it means to *implement* or *realize* a functional architecture (see Scheutz, 2000a).

2.2 Implementation of a Functional Architecture

So what are the implementation conditions for a functional architecture? To say that a system implements a functionalist description is to require that in addition to the input and output mapping, it has to get the mapping of the inner states right. Usually, these “inner

³ Of course, the behavior elicited by the organism realizing the cognitive system is specified as well.

states” are assumed to be multiply realizable, therefore the mapping has to be a many-to-one mapping from physical states to functional states (very much in the spirit of Chalmers, 1997). Yet, inner states are viewed by functionalists as intrinsically relational states, being *mutually defined* by all states in the functional architecture.

To illustrate this interdependence, consider, for example, the following automaton, which has two inner states ‘E’ and ‘O’ standing for “even” and “odd”. Depending on whether the number of ‘1’'s that the automaton has seen so far is even or odd, it outputs either ‘a’ or ‘b’, respectively.




Figure 1: The even-odd transducer with two inner states.

A functionalist account (e.g., see Block, 1996) of what it means to be in state E would look like this:

Being in E =_{def} Being an x such that $\exists P \exists Q [x$ is in $P \wedge$ (if x is in P and receives input ‘1’, then it goes into Q and outputs ‘b’) \wedge (if x is in Q and gets input ‘1’, then it goes into P and outputs ‘a’)].⁴

Since it is only claimed that there has to be an arrangement of physical states that corresponds to the functional states in a way that preserves inputs and outputs as well as transitions between states, it is possible for one physical state to serve as the instantiation of more than one functional state (and vice versa). Therefore, the correspondence between physical and functional states is not necessarily that of a mapping between physical types and functional types (let alone a 1-1 mapping), but rather that of a relation that preserves state transitions. “Implementation of a functional architecture”, therefore, has to be viewed as some sort of “bisimilarity” between functional and physical architecture rather than some sort of *isomorphic* relation from a functionalist point of view.⁵ As a consequence, finding a relation between a given functional architecture and a set of physical states together

⁴ Note that the existential quantifiers could be viewed as ranging over properties or as picking out particular physical states of the system.

⁵ The notion of “bisimilarity” is defined as follows: let I and O be two finite sets (e.g., the sets of input and output states, respectively) and let $M_1 = \langle S_1, \rightarrow_1 \rangle$ and $M_2 = \langle S_2, \rightarrow_2 \rangle$ be two structures with domains S_1 and S_2 , respectively, where relation \rightarrow_1 is defined over $S_1 \times I \times S_1 \times O$ and relation \rightarrow_2 is defined over $S_2 \times I \times S_2 \times O$. These structures are then said to be *bisimilar* if there exists a non-empty relation R between S_1 to S_2 such that for all $s_1 \in S_1$, $s_2 \in S_2$, $i \in I$, and $o \in O$ the following two conditions hold: (1) if $R(s_1, s_2)$ and $(s_1, i) \rightarrow_1 (t_1, o)$, then $(s_2, i) \rightarrow_2 (t_2, o)$ and $R(t_1, t_2)$, and (2) if $R(s_1, s_2)$ and $(s_2, i) \rightarrow_2 (t_2, o)$, then $(s_1, i) \rightarrow_1 (t_1, o)$ and $R(t_1, t_2)$.

with their causal transitions will be an intractable problem for reasonably large sets of states.⁶ In other words, a mere functional specification of a cognitive architecture is not going to be of any help in designing a realizer.

3 Behavioral States

3.1 An Ethological Perspective

To overcome the difficulties of tying functional specifications to physical implementations, I suggest to consider work done in animal behavior research as a venture point. According to animal behaviorists (e.g., McFarland, 1981), animal behavior can be categorized in terms of

- (1) reflexes (i.e., rapid, involuntary responses to environmental stimuli)
- (2) taxes (i.e., responses orienting the animal towards or away from a stimulus)
- (3) fixed-action patterns (i.e., time-extended sequences of simple responses)

While (1) and (2) are solely connected to external stimulation, (3) can have a contributing “internal” component as well (fixed action patterns can be “motivated”; take, for example, the “egg-retrieving” behavior of the greyling goose, see Lorenz, 1981, or Lorenz and Leyhausen, 1973). All three kinds of behaviors can be combined in complex ways to form hierarchies of behaviors (see figure 2).

In these behavioral structures, behaviors form “competitive clusters”, in which behaviors are mutually exclusive (e.g., in figure 2 the “fighting behavior” is such a competitive cluster comprising the mutually exclusive behaviors “chasing”, “biting”, and “display”).

To make these ideas of behavioral hierarchies more concrete, I will introduce the notion of *behavioral state*, which roughly corresponds to what is indicated by a “circle” in figure 2. Putting it crudely, a behavioral state is a state an individual is in if it performs a particular behavior (e.g., such as “food handling” or “looking out for prey”).⁷ “Behavior” is meant to be understood in a *wide sense* to include behaviors that are

not necessarily observable from the outside alone (such as “memory recall” or “thinking”, in general). Hence behavioral states are not simply combined input-output states, but rather they are some sort of “inner states” of an organism, states in which the organism is if it performs a particular kind of behavior. Note, however, that nothing is implied or claimed about a particular physical correlate of a behavioral state—it might or might not exist (I will return to this issue later).

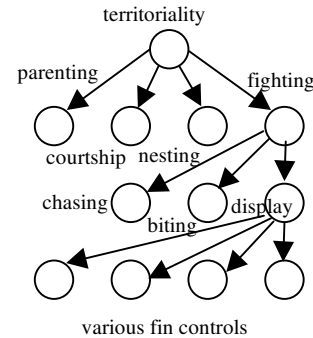


Figure 2: A part of a behavioral hierarchy for the male stickleback fish (see Lorenz, 1981). The various fin controls can be divided further into rays of each fin, the muscle fibers for each ray, and the motor neurons for each fiber.

Behavioral states are not restricted to “motor actions”, but include sensory actions as well as more abstract proprioceptive and reflective actions (such as monitoring inner physiological states, generating images, producing plans, recalling poems, analyzing pictures, making logical derivations, etc.). The latter ones are more “abstract behaviors”, which are mostly (if not completely) internalized and often involve solely parts of the cognitive architecture; in fact, they might not result in any externally observable change at all (a mathematician contemplating abstract objects and manipulating their representations in her mind might not need any stimulation from the outside world in performing this task, nor might any motor action result from it—this “brain in a vat”-idea with sustained cognitive activity whilst lacking external interaction seems to be at least conceivable in principle).

Memory and reflective processes, for example, are then viewed as special kinds of behavioral processes that lead to actions performed directly on the cognitive architecture, as opposed to the effectors of the individual.

In general, an individual will be in many behavioral states at the same time reflecting the fact that (1) some behaviors are contained in or shared among others (for example, searching for food as well as searching for a mate will both involve locomotion, despite the fact that the kind of search might be different), and (2) that many behaviors are performed in parallel (such as monitoring my hand as I move it to pick up an object).

⁶ This problem will indeed be at least as hard as “graph isomorphism”, which itself is believed to be in NP proper.

⁷ A note of terminology: while it is common usage to use “mental states” and “functional states” to refer to states of an individual’s mind, the notion of state is not exclusively used to describe “static” entities, but often times serves the role of a general term that subsumes *states* as well as *events*, i.e., processes. In a sense, the term “behavioral state” should have been avoided in favor of “behavioral processes”, as the latter emphasizes the dynamic character of the activity taking place in the individual. Following established terminology, however, I will continue using the term “behavioral state”, even if (systematic) dynamic changes in the individual are being referred to.

3.2 Behavioral Architectures

In a sense, the classical ethological picture outlined above is mainly concerned with the relation between various behaviors, it only depicts some causal relations between behaviors, and is, therefore, a functional specification of the behavioral architecture. Yet, partly implicit in and partly external to this picture is information about the time constraints as well as the strength of interactions and influences among behaviors (as studied and gathered by animal behaviorists). In other words, the picture is *incomplete* in so far as it leaves out essential implementation details that cannot be retrieved from a picture like figure 1 alone. Without these implementation details, however, some behaviors would not be the behaviors they are, since what distinguishes them from other behaviors might just be constraints on timing and strength of response (take, for example, a retraction reflex caused by touching a hot plate with your finger as opposed to the same movement being performed very slowly). Furthermore, the strength and configuration of interactions between behaviors is an integral part of their defining characteristics, which cannot be captured by a causal structure alone: suppose behavior A *causes* behavior B. Then this can happen in many different behavioral arrangements, for example, by A enforcing B directly or A suppressing C, which in turn inhibits B, or by A enforcing D, which enforces C, etc. Implicit in A (as defined by an animal behaviorist, say) is already information, which of these possible arrangements are realized in the animal. Hence, the causal structure might get restricted by the behavioral structure if (some of) the information implicit in the definition of behaviors is made explicit. In the following, I will briefly sketch how behavioral states can be defined to explicitly incorporate some of the otherwise implicit aspects of behaviors.

3.3 The Structure of Behavioral States and Networks

First and foremost, each behavioral state has an *activation level* and a *behavior* associated with it. This activation depends on various factors: (1) its own activation level, (2) the activation level of other states, (3) possible inputs from exteroceptive and proprioceptive sensors, (4) energy constraints, and (5) decay over time.⁸ The behavior associated with a behavioral states can either be a simple behavior (such as reflexes and taxes), or either a more complex fixed behavior (such as fixed action patterns) or a more complex adaptive behavior (which results from the interplay of fixed action patterns, reflexes, and taxes). The term “adaptive”

⁸ I will not be able to address issues related to last two points in this paper.

indicates that the latter kinds of behaviors can change over time, i.e., they can be learned, altered, etc. (utilizing the dynamic interplay of behavioral states).

Behavioral states are connected via inhibitory and excitatory links to other behavioral states and possibly to sensors (via “information channels”, i.e., filtering mechanisms that select parts of one or more sensory inputs and combine them in particular task-specific ways). Connections between behavioral states have a distance associated with them (expressed in terms of a time-lag), reflecting the “distance in space” that a signal has to travel from one locus of action to interact with another, allowing temporal as well as spatial integration of incoming signals.

Groups of behavioral states that are connected via mutually inhibitory links form so-called “competitive clusters”. They inhibit each other to various degrees, while usually entertaining excitatory connections to lower and upper level states (and possibly to some behavioral states of other clusters at the same level as well). In such a cluster the behavior associated with the state with the highest activation is activated and all behaviors of the other states are suppressed.⁹ This way hierarchical structures similar to the one in figure 1 can be defined which reflect the relationship between behaviors and in part also the complexity of each behavior associated with the various states (the lowest levels corresponding to simple reflex-like, reactive behaviors—this level has been explored in great detail in behavior-based robotics, e.g., see Arkin, 1992, or Brooks, 1986).




Figure 3: A hierarchy of behavioral states viewed as a two-layered architecture consisting of a deliberative and a reactive layer. Links with arrows indicate excitatory connections, links with circles inhibitory ones. The behavioral units in the deliberative layer do not operate on effectors, but perform internal operations (such as

⁹ There is evidence that similar mechanisms are at work in animals that inhibit all behaviors with lower activation values, e.g. see Lorenz (1981).

memory lookups, symbolic combinations, etc.).

With respect to the spread of activation, networks of behavioral states are very similar to I(nteractive) A(ctivation) and C(ompetition) networks (e.g., see Rumelhart and McClelland, 1986). Therefore, results from connectionist research about effects such as “blocking”, “settling”, “oscillation”, “hysteresis”, and others (often) apply *mutatis mutandis* to behavioral networks as well. The essential difference between IAC networks and behavioral networks is that the behavior associated with a behavioral state could affect via environmental feedback the activation level of the very state itself as well as the activations of other states. For example, a behavioral node representing the “search for black objects in visual field”-behavior might initiate ocular motor commands that lead to the detection of a small black object by another node, which in turn inhibits the search node, thus decreasing its activation, which in a mere IAC network would have otherwise not decreased.

As already mentioned, not all behaviors will involve physical effectors; in fact, only low level behaviors will directly exert influence on them (these are behaviors that would normally be localized in what roboticists refer to as “reactive layer”). Higher level behavioral states will mostly operate on structures internal to the cognitive system (these states would be situated in the “deliberative layer”). For example, a “retrieve image of mother” node (assuming for a moment there is such a node), might initiate a search in long-term memory (possibly involving other behavioral states) for a particular image that is associated with the individual’s mother. Or a “project-hand-move-forward” node might initiate a “simulated” hand movement in an emulator circuit, which is used to plan motions, resulting in a change in the circuit and as a consequence in other behavioral nodes (such as “collision detectors” in the emulator circuit, etc.).¹⁰ A behavioral network divided into a layered structure consisting of a reactive and a deliberative layer is schematically depicted in figure 3.

There are special cases of behavioral states that do not have any behavior directly associated with them. Instead of initiating an action directly, they contribute to behaviors indirectly by influencing other behavioral states, and can, therefore, assume the role of affective states. A state corresponding to “hunger”, for example, might receive inputs from proprioceptive sensors (i.e., a sensor monitoring the blood sugar or, more generally, the energy level) and exert positive influence on other states such as “search-for-food” (e.g., see Scheutz, 2000b). That way it is possible to entertain states that do not directly and immediately “cause” the individual

to act in a particular way, but might have indirect, long-term effects on the individual (e.g., depression, memory loss, etc.).¹¹

4 The Case for an Intermediate Level

4.1 The Relations between Physical, Functional and Behavioral States

So far, I have not explicated how physical and functional states relate to behavioural states as defined above. From an implementation perspective, behavioral states can be realized in many ways in different physical substrates. In brains, for example, they could correspond to a single neuron or to a group of neurons. They could be realized solely neuronally or maybe by involving other systems (such as the hormonal system) as well. Another physical medium, in which behavioral states can be realized, is the silicone of computers: computers can implement behavioral states by virtue of computational processes.

Some behavioral states might be (directly) “implemented” in the system in the sense that there exists a corresponding physical state or a set/sequence of physical states that are in *type correspondence* with the behavioral state. Other behavioral states might “supervene” on physical states in that there does not exist such a type correspondence—note that programs running on modern operating systems with virtual memory architectures exhibit such supervenience relations: when a program does not entirely fit into physical memory, it is loaded in parts on an “as-needed” basis, where different virtual memory locations get mapped onto the same physical memory location.

Another possibility for behavioral states to have no *fixed correlate* at all is to be only *partially* implemented (see Sloman, 1998) or to depend on environmental conditions (e.g., in terms of other behavioral states and/or environmental states—an example might be my performing the multiplication algorithm using paper and pencil: I am in a behavioral state which is implemented by a number of other states such as states of the paper and pencil, several visual routines, rule-retrieving memory processes and rule-following routines, etc.).

Behavioral states implemented in (sequences of) physical states are tightly coupled to their physical realizers (still allowing for multiple realizations), while behavioral states supervening on physical states do not exhibit such a coupling at all. They are realized by

¹⁰ I am currently investigating various possibilities of implementing simple emulator circuits in terms of behavioral states.

¹¹ Compare this to standard philosophical talk about “pain *causing* wincing and groaning, etc.,” where it is never clear whether pain always causes all the behaviors, exactly when the effects surface, whether showing the effects is necessary and/or sufficient for the individual to have pain, etc.

some physical states, but they might not show any systematic correlation to their realizers. For example, consider two networks of behavioral states, which are *functionally* identical except for the fact that the first explicitly implements a higher level behavioral state called “avoid-obstacle”, which is active if the agent is engaged in obstacle avoidance behavior. The second one does not have such a state, but can still control the same obstacle-avoidance behavior. In this case, the behavioral state “obstacle-avoidance” has a physical correlate in the former and no fixed physical correlate in the latter (what corresponds physically to the “obstacle-avoidance” state in the latter is a complex sequence of patterns that might, under different circumstances, not correspond to this state at all, e.g., if the agent follows another agent, which is avoiding obstacles, and thus is a “follow other agent” state, which by pure chance causes it to go through the same sequence of physical states... see also Pfeiffer and Scheier, 1999, ch. 12 for another example).¹²

This aspect of behavioral states seems very similar to the kinds of functional states about which philosophers tend to worry, and maybe most of the “high-level” functional states such as “belief states”, etc. are not directly (i.e., physically) implemented in the system (often the term “emergent” is used in this context). Even so, these kinds of behavioral states still retain one aspect lost in the mere “causation talk” of functional architectures, and that is *time*!

4.2 Causation and Time

It has been pointed out by philosophers (e.g., see Chalmers, 1997) that there is an essential difference between functional descriptions of physical systems like clocks, combustion engines, CD players, etc. and the functionalist descriptions of minds: in the former case some aspects of the physical structure matter, they are essential to any system realizing the functional architecture. Thus, these physical aspects are (if not explicitly, so then implicitly) retained in the functional architecture, thereby constraining the set of possible realizers. In the latter case, however, it is the very functional structure itself—so it is claimed—that matters, that is, the patterns of causal organization regardless of the underlying physical structure. Therefore, only causal organization, or put differently, “the flow of causation” is retained in functionalist abstractions from the physical as *the essential aspect* of minds. But is this really true?

Real minds are intrinsically tied to their environments and thus affected by the temporal structures imposed on them. Timing plays a crucial role in every

¹² Note that it should be possible to derive, beyond the causal properties, the temporal properties of the “obstacle-avoidance” state from the interaction of the (physically) implemented states.

aspect of a cognitive architecture pertaining to the proper functioning and survival of the organism. Many recent studies in cognitive science emphasize the importance of time as opposed to “mere temporal order” (see, for example, Port and van Gelder, 1995).

What distinguishes *time* from mere (temporal) *order* (as implicitly provided by the notion of causality) is that in addition to order a *metric* is defined (on the set of time points), that is, a notion of *distance* in time. This notion of distance in time allows one to differentiate functions according to their temporal behavior that would otherwise be indistinguishable. Take, for example, two microprocessors that work at different clock speeds—functionally they are identical, yet there is an essential difference between them, which is usually also reflected on any price tag put on them: their speed (another example of a function, where time is the distinctive factor, would be vowel production and recognition).

Is it problematic that causation alone does not suffice to capture the temporal structure of cognitive architectures? I would claim: Yes. Imagine two different physical systems that share the same functional specification of a human mind, one a regular human, another the People’s Republic of China “implementing the human brain” at a much, much slower pace (to use Block’s example). A human body controlled by the People’s Republic of China would fail terribly in the real world, because it could not react to its environment in due time.¹³ Well, one might say, it would do just fine if everything surrounding it, that is, its environment had been “slowed down” appropriately. This objection, however, strikes me as severely flawed, since it would entail *a completely new physics* (as in our physical universe certain processes have to happen at a certain speed otherwise they would not be the kinds of processes they are). Whether a “slowed down version” of a human mind could control a “slowed down version” in such a “slowed down universe” (with possibly completely different physical properties) seems too speculative a question to be taken seriously. What seems to be a productive approach, however, is to ask whether it is possible to understand a certain architecture (that evolved or was designed to meet the temporal constraints of its environment) at a mere causal level? I suspect that the answer would be *no* for systems that are sufficiently complex (like brains of vertebrates or VLSI microchips, for that matter).

If, on the other hand, causal structure were augmented by temporal constraints (i.e., information about

¹³ Many parts of our cognitive system have especially developed to meet time constraints of the environment. There is evidence for neural as well as chemical internal clocks (that work at certain clock rates), oscillator circuits that adapt to external cycles, etc. None of this would work if the system ran at 1/10000th of its regular speed. The same is true for digital circuits that have been designed to work at certain clock rates.

distance in time between causally connected states), then this would in theory suffice to capture an essential aspect of possible physical that implementations of the functional architecture. It would, for example, allow us to model the functional architecture computationally, i.e., to implement a virtual machine that abides to the temporal constraints (as many computational descriptions can handle temporal metrics, just take programming languages for real-time systems).

Behavioral states, therefore, seem to be an abstraction, which can be implemented computationally, and thus realized physically on computational systems. At the same time, behavioral states are abstract enough to capture aspects of minds that seem to be intrinsically connected to their causal structure and not to their physical realization (“organizational invariants” as Chalmers, 1997, puts it), thereby connecting them to functional descriptions of cognitive architectures.

5 Conclusion

The level of description of behavioral states is *intermediate* and *intermediary*, because it specifies states that could be realized in many different physical ways (in neural architectures, but possibly also in digital ones, and others), yet retains at least one crucial physical and causal aspect not retained in functional states: time! By explicitly incorporating time and thus allowing for modeling the temporally extended interactions between different states, this level might not only prove useful for constructing systems that exhibit complex causal interactions (such as minds), but also for explaining *how* functional states are related to physical states by viewing them as (not necessarily disjoint) collections of behavioral states.

References

- Arkin, R. C. Motor Schema-Based Mobile Robot Navigation. *International Journal of Robotic Research*, Vol. 8, No. 4: 92-112, 1989
- Block, N. What is Functionalism? *The Encyclopedia of Philosophy Supplement*, Macmillan, 1996
- Brooks, R. A Robust Layered Control System for a Mobile Robot. *IEEE Journal of Robotics and Automation*, Vol. RA-2, 1: 14-23, 1986
- Chalmers, D. J. A computational Foundation for the Study of Cognition, 1997 (published on the internet)
- Kim, J. *Philosophy of mind*, Westview, 1996
- Lorenz, K. *The foundations of ethology*, New York, Springer Verlag, 1981
- Lorenz, K. and Leyhausen, P. *Motivation and Animal Behavior: An Ethological View*, Van Nostrand Co., New York, 1973
- McFarland, D. *The Oxford Companion to Animal Behavior*, Oxford University Press, 1981
- Port, R. and van Gelder, T. *Mind as Motion: Explorations in the Dynamics of Cognition*, MIT Press, Cambridge, 1995
- McClelland, J. L. and Rumelhart, D. E. *Parallel Distributed Processing*, Vol. 1 and 2, MIT Press, Cambridge, 1986
- Pfeiffer, R. and Scheier, Ch. *Understanding Intelligence*, MIT Press, Cambridge, 1999
- Scheutz, M. Implementing Functional Architectures? 2000a (Submitted to VI. Congress of the Philosophical Society)
- Scheutz, M. Surviving in a Hostile Multi-Agent Environment: How Simple Affective States Can Aid in the Competition for Resources, *Proceedings of the Thirteenth Canadian Conference on Artificial Intelligence*, Springer Verlag, 2000b
- Sloman, A. Supervenience and Implementation: Virtual and Physical Machines, 1998 (submitted to ECAI)