

Reactive Robots and Amnesics : A Comparative Study in Memoryless Behavior

Amitabha Mukerjee

Center for Robotics
I.I.T. Kanpur, India 208016
e-mail: amit@iitk.ernet.in

Amol D. Mali (Member, IEEE)

Dept. of Computer Science & Engg.
Arizona State University,
Tempe, AZ 85287-5406
e-mail: amol.mali@asu.edu

Submitted - April 97, Accepted - September 97, Revised - August 98

Abstract

(Appears in **IEEE transactions on Systems, Man, and Cybernetics (Part C: Applications and Reviews)**, Volume 29, Number 2, May 1999, pp. 216-226.)

Reactive robot behavior has been claimed as a model of memoryless intelligence. Using recent results highlighting the cyclic conflicts in such systems, we throw light on similar phenomena in biology such as human amnesic behavior. In this paper, we present a formal characterization of cycles in the temporal sequence of memoryless behaviors and show that attempts to eliminate cycles result in less powerful behavior systems. Conversely, more powerful behavior systems are more likely to exhibit cyclicity. Such systems can be viable in performing simple tasks only by using the world as an external memory. Our results also provide useful insights into a better design of strategies for the treatment of amnesics.

“These memoryless robots will exhibit certain behavioral maladjustments, since, having no memory, they cannot react to events occurring over an interval of time. What reaction is suitable to a given input may depend on previous inputs, and such distinctions are a closed book to the memoryless robot. Also, these robots might easily get hung up or trapped even in a completely inanimate environment.” - James Culbertson, 1963 [10].

“Amnesics are conspicuously inaccurate in recalling the order of past events, and at times quite astoundingly imperceptive of the recurrences of similar occasions.” - George Talland, 1965 [32].

1 Introduction

Artificial intelligence (AI) has long owned up its debt to biology, but the gains to the biological sciences from the synthetic approaches of AI have been less prominent. Despite the differences in structure between brains and computers, the increasing similarities in functional objectives and methodology are making cross-disciplinary comparisons with synthetic intelligence models more and more relevant, even though these models lay no direct claim of cognitive or neuropsychological validity.

In this paper, we study a specific phenomenon that is exhibited similarly in memoryless systems in machines, animals, and humans – patterns of repeated behavior or cycles. This work is largely interdisciplinary, but is presented more from the Artificial Intelligence point of view than that of a cognitive psychologist. We consider a simplified memoryless, reflexive mode of behavior exhibited by a large class of reactive models used to design robot behavior. We conduct a formal analysis of such systems and relate our results to similar observations found in neuropsychological studies.

Particularly significant is the case where the consequences of one or more behaviors contain the stimulus for some previous behavior, resulting in cycles that can be potentially endless. The key result of this work is the finding that in memoryless systems, the likelihood of developing such cycles increases as the tasks attempted increase in complexity. We discuss existing mechanisms of breaking cycles like inhibition and claim that they do not reliably eliminate cycles. We also investigate some behavior modification mechanisms that are capable of reliably removing such cycles, but show how these result in a degraded performance.

Both artificial intelligence and biological intelligence use the term “behavior” to refer to some type of stimulus-driven actions. However the wide use of this convenient term has overloaded it with many differences of meaning, from the very complex human interactions characterized by Skinner, the specific patterns of ritual activity exhibited by animals (Lorenz), the input-output relation modulated by perceptual signals as in the motor schema model (Schmidt) to the boolean logic stimulus-response character of machine behaviors as in Brooks’ subsumption architecture [7]. Before we can create any formal models of behavior, it is imperative that we state exactly what is it that we are characterizing. We shall do this in the next section but first, let us introduce the notion of “memoryless behaviors” from the biological and artificial systems. In our discussion, biological memoryless systems are represented by amnesics and artificial memoryless systems are represented by the subsumption architecture-based robots of Brooks.

The relation between environmental change and action has sought to be related to specific characteristics of brain physiology. One of the richest lodes for this study has been the study of abnormalities or some form of brain deficiency. Neuropsychological research in past century has shown that damage to specific areas of the brain deprives a person of previously acquired knowledge and skill. Studies of impairment to memory owe much to the pioneering work by Korsakov. An interesting repetitive behavior was recorded by Korsakov for the microcephalic patient who repeatedly restarted dinner whenever food was brought in, without any recollection. This form of Korsakov syndrome (also called *amnestic confabulatory syndrome* as mentioned in [6]) results in the loss of memory of recent events, while impressions of long ago are recalled properly, so that the patient’s

personality is largely unaffected. Korsakov amnesics show near-normal effects of memory in their performance of a stem-completion task, even though their performance on direct tests of memory is severely impaired [16]. Eccles[11] reports an amnesic syndrome which causes a complete failure to establish memories for any happenings or experiences after the occurrence of bilateral destruction of the hippocampus and/or associated structures. This book also reports a patient with this amnesic syndrome who tries the same jigsaw puzzle day after day without showing any practice effect and reads the same magazines again and again without finding their contents similar. This forgetfulness also applies to people he has met since the lesion and even to those neighbors who have been visiting his house regularly for past six years. The patient has not learned their names and does not recognize them if he meets them in street.

“Could creatures who never remembered anything they perceived or experienced, have anything like human powers? Surely not” [19]. Several cases of amnesics are presented in [20] [32]. A Korsakov syndrome patient Jimmie thinks that he is thirty years younger and has no idea of the current date and when asked which month it is, he tries to look out and guess from the weather, using the world as an external memory [29]. He does not remember that he has talked to his doctor recently and greets him afresh as if meeting him for the first time [29]. Another type of cyclicity is exhibited in bodily functions - Jimmie starts to eat whenever food is presented, although he may have just finished his meal and may realize shortly thereafter that he is unable to eat. These types of cyclicity constitute a behavior parallel to that of memoryless robots that execute the same behaviors over and over again.

Animals with memory deficiencies also exhibit many instances of repeated cyclic behavior. Monkeys with medial temporal lesions constitute an animal model of human amnesia [38]. Repetitive behavior has been reported [25] in rats with dorsal hippocampal lesions. The running and sniffing in corners exhibited by a normal rat in a strange cage drops to a lower level after 10-15 minutes, but rats with damage to the dorsal hippocampus continue for at least an hour and probably for more than 12 hours. Kimble [17] reports that rats with hippocampal lesions made many repetitive errors in mazes. Passingham [27] reports that in the task of searching for food behind doors, operated monkeys (sulcus principalis group) returned fruitlessly to locations from which they had already removed the food, exhibiting a cyclic behavior.

In constructing synthetic models of intelligence, a roboticist is forced to simplify and abstract the biological systems, since biological systems of intelligence carry much evolutionary baggage that may not be relevant for their silicon-based counterparts [2]. Nonetheless, the construction of artificial systems is inspired by the study of natural animals. Beer *et al* [3], in their work on computational neuroethology, develop artificial nervous systems for insects to understand the neural mechanisms underlying the behavior of animals in their natural environments.

The neuroanatomist Braitenberg in his classic work *Vehicles* [5] describes a series of imaginary wheeled carts fitted with sensors and motors reflecting the idea of complex behavior arising from simple components. This paradigm of “synthetic psychology” has now been adopted by a large group of roboticists, who find it easier to model and debug smaller *behavior modules* as opposed to larger and more integrated centralized controllers.

Interesting results have been achieved using this strategy in the can collection robot [9], navigation of mobile robot [2], office rearrangement tasks in robot competitions, etc. The behavior-based intelligence paradigm propounded by Brooks and others has challenged the role of representation in AI. Brooks argues that when intelligence is approached in such an incremental manner, the reliance on representation disappears [8].

In response, traditional AI researchers and cognitive scientists such as Kirsh have argued that control cannot serve as a complete substitute for representation [18]. Vera & Simon [36] also show that internal representations are necessary for memoryless artificial systems if they are to perform really complex tasks. Boden [4] reports that even simple animals form representations and that experiments show that monkeys are sensitive to temporal order. At the same time, behavior systems have also been moving away from the memoryless or purely reactive paradigm. A well-known extension of the Brooks' approach includes building maps [7] in a module that learns what looks suspiciously like a central representation. Our results however, are limited to the memoryless, purely reactive systems.

Recently, cyclic behaviors similar to those of amnesics, have begun to show up in artificial reactive systems. Connell records an instance where a can collecting robot attempts to re-pick the can it has just deposited in the destination area as shown (Figure 1); this conflict was detected only after a full implementation [9]. Cyclical wandering and cyclic conflict of going back and forth between two obstacles have been reported by Anderson & Donath [1]. Even the very early reactive systems such as the "tortoises" constructed by Grey-Walter in the 1960's could find themselves in cycles. Anderson & Donath [1] say that

such cyclical behavior may indicate that an essential characteristic of a truly autonomous robot is that of having and using memory (i.e., the ability to react to events which occur over a number of intervals of time and the ability to alter behavior depending upon previous behaviors). Our results put this intuition on a more formal basis by showing that the absence of temporal global memory in particular leads to a lower modularity (defined in 2.3) in behavior systems. Figure 2 shows a general model for cyclic conflicts where agents A, B, and C stimulate each other in an endless cycle.




Figure 1: *Conflict in can picking.* After dropping the can, the stimulus for the *pickup* behavior is immediately activated if the robot is looking towards its dropoff point.




Figure 2: *Cyclic Conflict.* Agent A’s action triggers Agent B. B triggers C, which calls for action by A again, endlessly.

Our analysis in this paper is dependent on a crucial observation regarding the temporal structure of memoryless systems. The conflicts we are addressing are temporal sequence conflicts for which it is necessary to define the temporal structure of behaviors, which is usually sequential since one behavior usually provides the stimulus for another, so that

there is often a clear temporal sequence in which behaviors are executed. In this paper we show that cycles occurring in this temporal sequence can be avoided reliably only by modifying the behaviors themselves, and we introduce two such modifications, based on specializing the stimulus or restricting the action of a behavior. One of the key results of the paper is that any such modification reduces the usefulness of the behavior structure and makes it less flexible. The paper is organized as follows - in section 2, we develop a formal model of reactive systems and use it to prove our results in section 4. In section 3, a classification of behavior conflicts is provided and the potential of inhibition to eliminate cycles is discussed. In section 4, we discuss our behavior refinement mechanisms for reliably eliminating cycles and prove our results on the effects of behavior refinements and modularity of behavior systems. In section 5, we discuss the insights we obtained on the connection between synthetic psychology and neuropsychology. Section 6 presents our conclusions.

2 What Is a Behavior?

Ethologists, psychologists and cognitive scientists are mostly interested in analyzing pre-existing behaviors, but synthetic biologists and roboticists are interested in creating models of behavior. Each of these groups uses the term behavior in senses that are related but are fundamentally different.

Brooks looks upon behaviors as types of an intelligent module, an input-output relation to solve small problems [7]. Hopefully these modules can be combined to solve larger problems. There is no shared global memory. The stimulus to a behavior is boolean and

is tested by an *applicability predicate* as opposed to the continuum of activation (where the strength of a stimulus matters, rather than just its presence) typically noted in biological behaviors. Each behavior intermediates directly between the input stimulus and the output action, and the inter-process communication is kept at a minimum [7]. In this paper, we investigate this robotic model of behavior in principle, and show how some of the results we establish also apply in biological domains.

There also exist other models of artificial behavior. Minsky suggests thinking about a goal directed behavior as an output of a difference engine that measures the differences between the current world state and the goal state and takes actions to reduce these differences [26]. This has the flavor of means ends analysis. Tharp *et al* conclude that new and complex behaviors rarely emerge spontaneously in their complete form, rather than awaiting such an improbability, a wise interventionist will begin the teaching process from that point in the behavioral repertoire which is the closest, if still imperfect approximation to the ultimate criterion [33]. This sounds like planning. Arkin proposes the *perceptual and motor schemas* as a model of behavior specification for the navigation of a mobile robot [2]. Next, we develop our notation.

2.1 Notation

In this work we have followed the behaviorists and adopted a *3-tuple* model of behavior: $\langle stimulus, action, consequence \rangle$. A behavior module β takes the form $\langle s, a, c \rangle$. However, the action a is not directly referred to by us, and we sometimes abbreviate the notation to $\langle s, c \rangle$. The behaviors we refer to have a primitive action component that is directly executable. Hence building a house is not a behavior as per our definition,

rather it is an abstract task that may be fulfilled if the behaviors get chained properly. In the examples and proofs that follow, the stimulus s and the consequence c are defined in terms of a conjunction of predicates as in Brooks' subsumption architecture. However, any type of activation, where the consequence of one behavior leads in turn to the stimulus for another, could be substituted, and the basic structure of our analysis and the results would remain unchanged. The following symbols are standard to predicate calculus: \wedge - conjunction, \vee - disjunction, \Rightarrow - implication, \sim - negation, \forall - universal quantifier, \exists - existential quantifier. We note that any behavior has two states - *waiting-for-stimuli*, which may be called the dormant state, and the actual action, which may be called the dominant state. We define the *dominant period* of a behavior as that period when the behavior is active. In most behavior implementations, behaviors become dominant in a temporal sequence. We use the symbol “:” (*precedes*) to denote this. $\beta_1 : \beta_2$ implies that behavior β_2 becomes dominant following the behavior β_1 .

2.2 Behavior Chain

We define a *behavior chain* as a sequence of behavior modules $\{\beta_1 : \beta_2 : \beta_3 : \dots : \beta_n\}$. Here the action of the earlier module changes the situation in such a way that the newly changed part of the situation is in turn a stimulus for the next module in the sequence. If the consequence and stimulus include a finite universal state as well, then we can say that the stimulus s_{i+1} of the behavior module β_{i+1} is logically implied by the consequence of the module β_i i.e. $(c_i \Rightarrow s_{i+1})$. What we mean by the finite universal state can be clarified by an example. Let $Universe = X \wedge Y \wedge Z$ and $c_1 = A$ and $s_2 = X \wedge A$. Then β_1 leads to β_2 but $(c_1 \not\Rightarrow s_2)$. This is another version of the frame problem, which arises in the add and

delete list type notation as well [12] (Some planning systems that use operators represent consequence of an operator in terms of an add list and a delete list, i.e. they separately list predicates that are made true and false). One of the methods for overcoming this is to have the state of the *universe* included in all stimuli and consequences in the system. Thus when we say that $(c_1 \Rightarrow s_2)$, we mean that a part of s_2 was true in the *universe* and some literals in c_1 cause the rest of s_2 to come true. Thus in order for $(c_1 \Rightarrow s_2)$ to be true, both the stimulus and consequence should always contain the state of the *universe*. This allows us to develop the argument effectively, skirting the philosophical debate on the frame problem.

Such a chain of behaviors is observed in biological systems too, e.g. Tinbergen [34] reports that a stickleback's courtship behavior can be characterized by the chain - Female appears, Female gives head-up display, Male swims zigzag to female, Female swims head up towards male, Male swims towards nest, Female follows the male, Male shows the nest, Female enters the nest, Male tremble-thrusts, Female spawns, Female leaves, Male enters, Male fertilizes.

We define a *behavior space* B as a set of behavior modules. A temporal chain of behaviors C is said to be composable from B (written as $C \triangleleft B$), if and only if $(C = \text{ordered set } \{\beta_i\} \wedge (\forall \beta_i \in C) (\beta_i \in B))$. A *stimulus space* Σ of a behavior space B is the union of the stimuli of all behavior modules in B .

2.3 Power, Usefulness and Flexibility of Behaviors

To compare different behavior systems, we define a few relative measures.

Definition 1

• **Power** : A behavior $(\beta := \langle s, a, c \rangle)$ is more powerful than $(\beta' := \langle s', a', c' \rangle)$ iff $(s' \Rightarrow s) \wedge (c \Rightarrow c') \wedge (\beta \neq \beta')$.

In other words, the more powerful behavior can be triggered at least as frequently as the less powerful one and results in at least as strong a consequence. In Figure 3, behavior C is most powerful since it has the most general stimulus, though A, B and C all have similar consequences.

This argument can be generalized for the entire behavior space. A behavior space B is more powerful than the behavior space B' if B' can be obtained from B by replacing some behavior(s) $\beta \in B$ by less powerful behavior(s) β' .


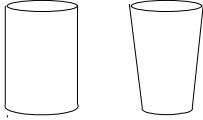
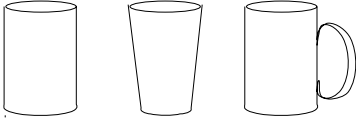
Stimuli for behaviors of increasing power		
(A)	(B)	(C)
		
$graspable(x) \wedge can(x)$	$graspable(x) \wedge (can(x) \vee cup(x))$	$graspable(x)$

Figure 3: *Power*: Behavior B can pick up cups as well as cans (has a more general stimulus). C is even more general. The notion of power is used to capture this, provided all three behaviors have the same consequence.

• **Usefulness** : We define a *task* as a transition from one state of world to another, achieved by a temporal chain of behaviors. A behavior space B spans the task space τ if and only if $\forall (t \in \tau) (\exists (C \triangleleft B) fulfills(C, t))$. The *greatest potential task space* $\tau_G(B)$ is the largest task space that is spanned by the behavior space B . The *usefulness* of a behavior space is defined as the ratio $\frac{|\tau_G(B)|}{|B|}$.

• **Flexibility:** A behavior space B is at least as flexible as the behavior space B' if $(\forall t \in (\tau_G(B) \cap \tau_G(B'))) \exists (C \triangleleft B) (\text{fulfills}(C, t) \wedge (\forall (C' \triangleleft B') (\text{fulfills}(C', t) \Rightarrow |C| \leq |C'|))$). This means that B can fulfill tasks with shorter chains (which are composed from fewer behaviors). It is possible that $\tau_G(B_1) = \tau_G(B_2)$ but $|B_1| \ll |B_2|$ because the behaviors in B_1 allow a richer composition of chains that are shorter. The flexibility metric captures this. Here higher flexibility is also correlated with higher usefulness. But note that the semantics of flexibility and usefulness is considerably different. Consider the case where $|B_1| = |B_2|$ and $\tau_G(B_1) = \tau_G(B_2)$ and hence these behavior spaces have the same usefulness. However the user of B_1, B_2 may be interested in only a proper subset of $\tau_G(B_1)$. Hence even though both B_1, B_2 have the same greatest potential task space and usefulness, B_1 may be more flexible and preferred by the user over B_2 .

• **Modularity:** A behavior space is more modular if different modules in that space are more independent, in the sense of minimal interference between its modules. One measure of interference in a behavior space is the incidence of cyclic behavior. We therefore define the modularity of behavior space B as an inverse of the likelihood that cyclic conflicts will arise in B . The properties of these metrics are discussed in [24].

3 Detection of Conflicts

In the broad sense of the word conflict, any behavior chain leading to non-fulfillment of the desired objectives can be said to have a conflict. Let the chain $C = \{\beta_1 : \beta_2 : \dots : \beta_n\}$ be the desirable behavior sequence that achieves a desirable outcome. There are three types of conflicts that can prevent the chain C from being executed, by breaking the sequence

$\beta_i : \beta_{i+1}$.

Definition 2

(a) **Extraneous behavior Conflict**: $\beta_i : \beta', \beta' \notin C$.

(b) **Cyclic Conflict**: $\beta_i : \beta_k, \beta_k \in C, k \leq i$. (discussed later)

(c) **Skipping Conflict**: $\beta_i : \beta_k, \beta_k \in C, k > (i+1)$. This type of conflict can be treated in a manner analogous to extraneous behavior conflicts by prioritizing the behaviors.

The type of behavior that we are investigating is the *cyclic conflict*, where, both β_{i+1} and β_k may be triggered and clearly the triggering of β_k would lead to a cycle (Figure 4).




Figure 4: *Cycle in a temporal chain of behaviors.* Consequence c_i triggers the stimulus s_k of an earlier behavior. $\beta_i : \beta_k$ should be broken but the chain from behavior β_i to β_{i+1} should remain intact.

3.1 Terminating Cycles

All cycles do not result in a conflict. Let us say that a tower of specific height needs to be built with a number of small blocks. Let us say that the module *build_tower* repeatedly stacks the blocks one above the other until the tower has the required height. Although this involves a cycle, there is a clear termination condition implied by the height of the tower, hence this is not a cyclic conflict. One of the difficulties of robots and other synthetic behavior systems is distinguishing the fatal cycles from the desirable ones, a problem that

can be shown to be equivalent to the well known *halting* problem, and is known to become prohibitively expensive to solve as the number of modules grows.

3.2 Inhibition

One of the conflict handling methods widely observed in natural behavior systems is that of behavior inhibition, where one behavior inhibits another from performing its normal task. The inhibitive effect may or may not continue beyond the end of the dominant period of the inhibiting behavior. We define two types of prioritization based on this.

- *inhibition* or *suppression*, a module γ suppresses the output of β , and instead of $\alpha : \beta$, $\alpha : \gamma$ occurs, where α is the behavior module that was dominant immediately before the inhibition occurred.
- *Delayed action* is a special case of inhibition, where the inhibitive link remains effective for some time t_{delay} even after the inhibiting module is no longer dominant.

Connell uses the term *retriggerable monostable* to capture this sense [9].

These mechanisms are not guaranteed to kill the stimulus of β_k , hence β_k may be active after the dominant period of the suppressing module is over. Thus within the scope of the prioritization schemes discussed here, it is not possible to guarantee that cyclic conflicts in artificial systems will be avoided.

3.3 Detecting Cycles in Behavior Graphs

Mali & Mukerjee [23] mention that cyclic behavior can be detected by detecting cycles in the temporal graph constructed by connecting individual behavior chains. However this is subject to identifying all terminating cycles, which can be done, for example, by

flagging all such loops at the programming time. Similar models can also be constructed for memoryless biological systems. For example, in one of the examples in [27], in the food search task, the unoperated control group does not revisit nodes, but the sulcus principalis group of monkeys revisits some of the nodes several times, indicating the presence of some type of cycle in their behavior graph (Figure 5). Similar cycles have been observed in the reactive navigation of robots [2].




Figure 5: First 20 moves made by each animal on Trial 21 of the search task. (The star gives the first door opened, and each move is numbered. UC = unoperated control, SP = sulcus principalis.) (Figure reproduced from [27].)

4 Behavior Refinement

In this section, we investigate the very viability of memoryless systems. Can a system exhibiting cyclic conflict be modified in such a manner as to avoid these cycles? What are the effects of such modifications?

One of the most important types of modification of behaviors in biology is that of *Conditioning*, or behavior refinement, whereby a behavior adapts to a new environmental reaction. In artificial intelligence parlance, this is typically known as “learning.” Learning can be of two types; *adaptation* where a behavior is adapted to meet a new but similar environment; and *generalization* where a behavior is generalized to be able to handle both new and old environments. For example, the behavior *pick-up-cup*, may be adapted to become *pick-up-soda-can*, which would work with soda cans. A more powerful form of learning would result in the generalized *pick-up-drink*, which can pick up either a coffee cup or a soda can (Figure 3). In the latter instance, the stimulus has been generalized, and the consequence remains the same, so this new behavior is more powerful than the old one in the sense of the definition of *power* in section 2.3.

In biological memoryless systems however, the capability to learn may or may not survive, depending on the type of memory loss or amnesia. In addition, as we show, memoryless-ness leads to the higher likelihood of conflict for the learned, more powerful behaviors.

From definition 2, whenever there is a cycle in a behavior chain $C = \{\beta_1 : \beta_2 : \dots : \beta_k : \dots : \beta_i : \beta_{i+1} : \dots : \beta_n\}$, there must be a triggering of the kind $\beta_i : \beta_k$, where $k \leq i$. Then both $\beta_i : \beta_{i+1}$ and $\beta_i : \beta_k$ can occur at this point. Our objective is to break the

$\beta_i : \beta_k$ link without disturbing the $\beta_i : \beta_{i+1}$ or $\beta_{k-1} : \beta_k$ triggerings which are essential to the successful execution of the chain. We have seen that priority-based methods are not guaranteed to achieve this, so we look for behavior modification approaches which will maintain $(c_{k-1} \Rightarrow s_k)$, $(c_i \Rightarrow s_{i+1})$ and break $(c_i \Rightarrow s_k)$. We develop two methods for achieving this: in stimulus specialization, s_k is specialized, and in response generalization, c_i is generalized.

4.1 Stimulus Specialization

Let us consider the conflict in picking up the soda cans, where the freshly deposited can is picked up. If we were to add the condition “not-deposited-just-now (x)” to the stimulus for the behavior *pickup*, then we would need only a small recency memory (recently dropped can). Thus the stimulus for β_k becomes more specialized. However, in doing this, one must be careful so as not to disturb the rest of the chain, i.e. $(c_{k-1} \Rightarrow s_k)$ should still hold but $(c_i \Rightarrow s_k)$ must not. Clearly this will not be possible where $(c_i \Rightarrow c_{k-1})$, then any changes we make in s_k such that $\beta_i : \beta_k$ is broken will also result in $\beta_{k-1} : \beta_k$ being broken (which breaks the chain in an undesirable way). Thus stimulus specialization can be used only if $(c_i \Rightarrow c_{k-1})$ is not true. To break the cycle without disturbing the rest of the chain, consider the scenario where there exists a literal γ such that $(c_{k-1} \Rightarrow (s_k \wedge \gamma))$ but c_i does not imply $(s_k \wedge \gamma)$. We call the conjunction of all such literals $\Gamma = (\gamma_1 \wedge \gamma_2 \wedge \dots \wedge \gamma_m)$, the maximal difference between c_{k-1} and c_i . Stimulus specialization works only when $\Gamma \neq \emptyset$, and involves modifying s_k to $(s_k \wedge \gamma)$, $\gamma \in \Gamma$. It is advisable not to overspecialize s_k (e.g. by adding excessive literals), since this adversely affects the power of the behavior β_k . For the purpose of understanding, one can consider a simpler case where both c_i and c_{k-1}

are in purely conjunctive form. Then Γ is nothing but the difference $(c_{k-1} - c_i)$ and s_k is modified by conjunction with one of the literals that is in c_{k-1} but not in c_i . Note that since the stimulus is specialized, any stimuli that are required by the action are still available to it.

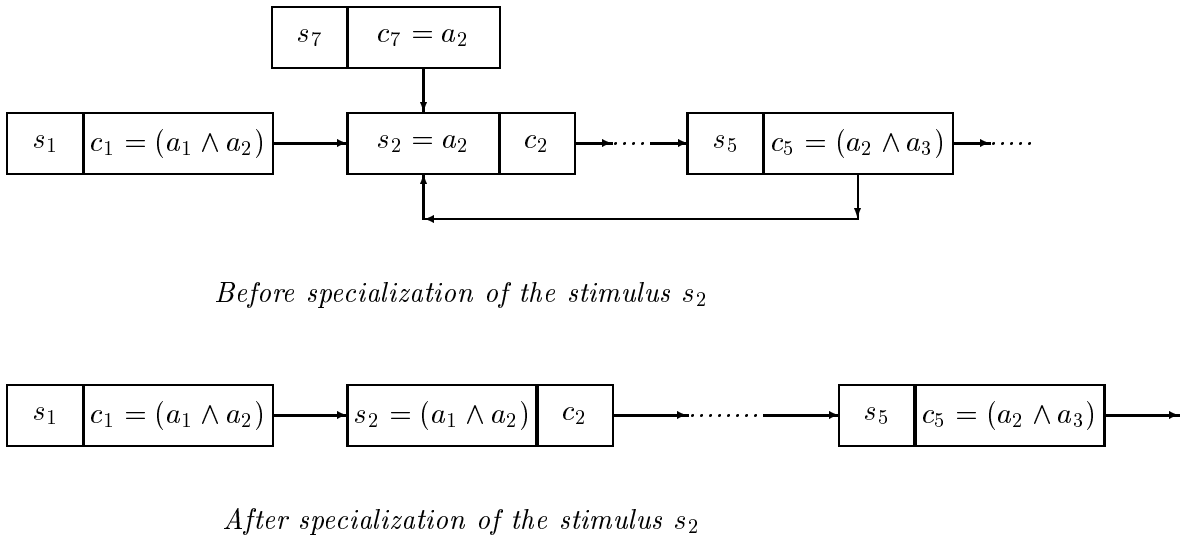




Figure 6: *Effect of stimulus specialization (Though the cycle is broken due to stimulus specialization, the chain $\{\beta_7 : \beta_2 : \dots : \beta_5 : \dots\}$ is no longer composable).*

4.2 Response Generalization

Here the action is modified so that the consequence of the action is weaker i.e. if the old consequence was c and the new one is c' then $(c \Rightarrow c')$ but c' does not imply c . For example, we can modify the action of the module *drop* so that while dropping the can on the ground, the robot puts it in an inverted position which prevents the robot from detecting that the object is a can. The original consequence was $\exists x (can(x) \wedge graspable(x))$ and the modified



Before generalization of c_5



After generalization of c_5

Figure 7: *Effect of response generalization (Though the cycle is broken due to response generalization, the chain $\{\beta_1 : \beta_2 : \dots : \beta_5 : \beta_7\}$ is no longer composable).*

consequence is $\exists x \textit{graspable}(x)$, assuming the sensors cannot identify cans that have been inverted. Otherwise, we may modify the consequence by covering the can to make the predicate $\textit{can}(x)$ false, then this leads to addition of a new behavior module or modifying the action part of the original module, both of which require considerable re-programming, and are expensive. In response generalization, $(c_i \Rightarrow s_k)$ must be negated, while $(c_i \Rightarrow s_{i+1})$ must hold. Hence response generalization can be used only when $(s_{i+1} \Rightarrow s_k)$ does not hold. In fact, one could characterize the process of response generalization by saying that there must exist a literal σ s.t. $c_i = (c'_i \wedge \sigma)$ and $(c'_i \Rightarrow s_{i+1})$ ($(c_i \Rightarrow s_k)$ holds from the definition of cycle) but c'_i does not imply s_k . Again, if s_k and s_{i+1} are in conjunctive form, then modifying c_i is better understood as dropping the literal σ already appearing in c_i , written as $(c_i - \sigma)$. Since stimuli/consequences are often conjunctive, this difference notion is a useful concept in practice. Thus c_i is modified to $(c_i - \sigma)$, where $\sigma \in (s_k - s_{i+1})$.

Response generalization requires that the action should be modified, whereas stimulus specialization only makes the same action more restrictive. However, stimulus specialization may not always be possible; e.g. with “not-deposited-just-now(x)” the robot may still pick up a can that was not dropped recently. Better solutions to this, such as “never-deposited-before(x)” or “not-at-repository(x)” would require considerable memory. Therefore, stimulus specialization, while cheaper to implement, may not be effective in many instances, since the right level of specialization may require large memory and specializing with limited memory may result in an undesired behavior.

Though this cycle elimination analysis has been developed for artificial systems, similar solutions may apply in the domain of biological memoryless systems like amnesics.

Here specializing the stimulus may be seen as creating a reminder message in the amnesics' environment. In case of Jimmie [29] who wants to talk to the same doctor again and again without remembering that he has recently had a conversation with his doctor, photographs of people with whom Jimmie had a conversation can be kept in the external world near Jimmie. There is some evidence of the ability of external sources of information to strengthen recollections in humans, lowering the effect of childhood amnesia (childhood amnesia refers to forgetting events in childhood and not being able to report them when one is an adult) as discussed in [35].

4.3 Effects Of Behavior Refinement

Let us now investigate the effects of stimulus specialization and response generalization.

Lemma 1(a). Whenever a behavior space B is modified to B' to eliminate cyclic conflicts by specializing stimulus s of some behavior $\beta \in B$ to s' so that $(s' \Rightarrow s)$, then $\frac{|\tau_G(B)|}{|B|} > \frac{|\tau_G(B')|}{|B'|}$ (the usefulness of the behavior space decreases).

Proof :- Let Σ be the stimulus space and $s \in \Sigma$ be a stimulus in conjunctive form. Now let s be specialized to s' so that $(s' \Rightarrow s)$. Now tasks or subsequent behaviors fulfillable or triggerable in $(s - s')$ (this difference corresponds to the states of the world in which s holds but s' does not) will no longer be so. Thus we need a new behavior β'' such that s'' will be triggered in those states of the world in which s is triggered but s' is not, so that β' and β'' together serve the stimulus set s which implies that $|B|$ increases and the usefulness of B decreases (as per the definition in 2.3) (We assume that $c'' = c' = c$ since our aim is to restore the functionality of the behavior space back to its original with as

few local changes as possible. However Lemma 1(a) also holds when either c' or c'' or both are weaker than c). If the new behavior is not added, the usefulness decreases because of decrease in the greatest potential task space. If there is another behavior $\beta''' \in B$ such that $s''' = s$ and $c''' \neq c$ (otherwise β and β''' will be the same), then will β''' maintain $|\tau_G(B)|$ constant? This cannot happen because either (i) just like s , s''' will also have to be specialized to eliminate cycles or (ii) β''' will not participate in the cycle and not keep $\tau_G(B)$ constant since it has a different consequence. Hence no other existing behavior in B can maintain $|\tau_G(B)|$ constant. Hence the size of $\tau_G(B)$ reduces. The argument in this proof can be generalized to non-conjunctive stimuli, where also a similar set of unserved stimuli can be found. \square

If a new behavior is added, length of the chain required to fulfill unfulfilled tasks will not necessarily increase, hence we do not claim that stimulus specialization does reduce flexibility. However this is not the case with response generalization. This leads us to the next result.

Lemma 1(b). Whenever a behavior space is modified to eliminate cyclic conflicts by response generalization, the usefulness of the behavior space decreases.

Proof :- Let the consequence c of β be generalized to c' so that $(c \Rightarrow c')$ keeping the stimulus same ($s = s'$). Thus $(c - c')$ is not being made available by β . Let the behavior space containing β' be denoted by B' (the spaces B, B' differ only because $\beta \in B$ is replaced by β' to derive B' from B , all other behaviors from these two spaces are the same). Hence new behaviors like β'' are needed to complete the tasks requiring $(c - c')$ and this change

increases $|B|$. This implies that the usefulness of B decreases (as per the definition in 2.3). (We assume that $s = s' = s''$ since our aim is to restore the functionality of the behavior space back to the original with as few local changes as possible. We also assume that β', β'' execute without any control conflict (which may arise due to contention for resources) to provide c such that $c = (c' \wedge c'')$. However Lemma 1(b) also holds when s' or s'' or both are stronger than s .) Also the addition of the new module or a set of modules increases the lengths of the chains composed to fulfill these tasks resulting in a decrease in the flexibility. If this not done, then some tasks requiring $(c - c')$ (this should be interpreted as a pure difference of literals) cannot be performed which implies that $|\tau_G(B')| < |\tau_G(B)|$ which means that the usefulness of the behavior space decreases.

If there is another behavior $\beta'' \in B$ such that $c'' = c$ and $s'' \neq s$ (otherwise β and β'' will be the same), then will β'' maintain $|\tau_G(B)|$ constant? This cannot happen because either (i) just like c , c'' will also have to be generalized to eliminate cycles or (ii) β'' will not participate in the cycle and not keep $\tau_G(B)$ constant since it has a different stimulus. Hence no other existing behavior in B can maintain $|\tau_G(B)|$ constant. Hence the size of $\tau_G(B)$ reduces.

The argument in this proof can be generalized to non-conjunctive consequences, where also a similar set of unserved stimuli can be found. \square

Let us say that we have a behavior β whose consequence $c = (p \wedge q)$ leads to a cycle. If we use response generalization, (and hence design action of the behavior in such a way that the consequence c after modification contains either p or q but not both), we may

have to design two behaviors β' and β'' such that $c' = q$ and $c'' = p$. If β has the stimulus $s = (p \vee q)$ which is triggered leading to a cycle and if we use stimulus specialization (which means dropping p or q from the current stimulus s), we may have to design two more behaviors β' and β'' such that $s' = p$ and $s'' = q$. If *pick_up* is specialized so that it does not pick up a can it has dropped, then to drink coffee, we will have to design another “multiple_pick_up”. In some cases, it may not be possible to design an action that will fulfill these conditions. This discussion brings us to our most important results, which have to do with the power and usefulness of behavior spaces.

Behavior Modification Theorem. Given two behavior spaces B and B' such that B is more powerful than B' (i.e. B' is obtained from B by replacing some behaviors β of B by the less powerful ones β') then:

(a) The greatest potential task space of behavior space B' is smaller than that of B , i.e.

$$|\tau_G(B')| < |\tau_G(B)|$$

(b) Usefulness of B is larger than that of B' i.e.

$$\frac{|\tau_G(B)|}{|B|} > \frac{|\tau_G(B')|}{|B'|}.$$

(c) Likelihood of a cycle in B is at least as large as that for B' .

Proof (a, b) :- First, let us consider the case where a single behavior β has been replaced by the less powerful β' . The set of chains of behaviors composable from a behavior space represents a tree with the root corresponding to the availability of the right initial stimulus and each node in this tree represents a world state which may correspond to the desired state, indicating an existence of a task fulfilling chain. The greatest potential task space is proportional to the total size of this tree of behavior chains (this size is defined in terms

of the number of all possible chains contained by the tree and the number of such trees, since different trees can be constructed for different initial world states. Such trees can be combined into a single tree, the root of which is a node that represents virtual initial state. This root can then be connected to the actual initial states (roots) of the individual trees by links). Now, either the behavior β will have more applicability due to a weaker stimulus as compared to the behavior β' , or the behavior β will have stronger consequence resulting in more behaviors being triggerable after it or both. In terms of the task tree, either β will have more parent nodes, or it will have more children. In either case, the branching factor in B is larger than that in B' and the size of the task tree will be larger. Hence the result (a). Since $|B|$ has not changed, the usefulness of the behavior space, $\frac{|\tau_G(B)|}{|B|}$ decreases when β is replaced by β' which proves part (b). This treatment can be extended to multiple instances of replacing a stronger behavior by a weaker one. \square

(This argument has a connection with biological systems too, at neural level, branching at axons and dendrites enables a single neuron to interconnect with a huge number of receptors, effectors and other neurons. This establishes the cellular basis for behavioral flexibility that characterizes animals with complex nervous systems [30]. In biological systems, this holds at behavioral level too. Stronger consequence enables more actions and behaviors with weaker stimuli have a wider applicability.)

Proof (c) :- Let $\beta_i \in B$ and $\beta'_i \in B'$ be two behaviors s.t. β_i is more powerful than β'_i , i.e. $(s'_i \Rightarrow s_i)$ or s_i is weaker than s'_i . The spaces $B - \{\beta_i\}$ and $B' - \{\beta'_i\}$ are the same. Now consider any chains of length n composable in B and B' , which differ only in that the module β_i is replaced by β'_i . Consider all behaviors $\beta_j \in C$, $j \geq i$, with consequence

c_j . The likelihood of a cycle in B , denoted by $L_{cycle}(B)$ (which we do not restrict to 0-1 range like probabilities) is

$$\sum_{j \geq i}^n (prob(c_j \Rightarrow s_i)),$$

and $L_{cycle}(B')$ is

$$\sum_{j \geq i}^n (prob(c_j \Rightarrow s'_i)).$$

Clearly, since $(s'_i \Rightarrow s_i)$, $\forall c_j (prob(c_j \Rightarrow s_i) \geq prob(c_j \Rightarrow s'_i))$. Similarly $(c_i \Rightarrow c'_i)$ for which similar analysis can be carried out. Thus $L_{cycle}(B) \geq L_{cycle}(B')$. If there is a cycle in B' , there will be a cycle in B too. \square

Modularity Theorem. Given cycle free behavior spaces B and B' (B' is obtained from B by replacing some behaviors of B by less powerful ones) and a module λ not belonging to B and B' is added to both of them, then the space $B \cup \{\lambda\}$ is at the most as modular as the space $B' \cup \{\lambda\}$.

Proof - The theorem follows from the part (c) of behavior modification theorem. As B contains more powerful behaviors, the chance of occurrence of a chain $C = \{\beta_1 : \beta_2 : \dots \beta_n\}$ such that $(c_i \Rightarrow s_k)$, $i \geq k, 1 \leq i, k \leq n$ is higher in $B \cup \{\lambda\}$ and from part (c) of behavior modification theorem, $L_{cycle}(B) \geq L_{cycle}(B')$. Thus, $L_{cycle}(B \cup \{\lambda\}) \geq L_{cycle}(B' \cup \{\lambda\})$. Hence $\frac{1}{L_{cycle}(B \cup \{\lambda\})} \leq \frac{1}{L_{cycle}(B' \cup \{\lambda\})}$. This decreases modularity as defined in section 2.3, hence the space $B \cup \{\lambda\}$ can be at the most as modular as the space $B' \cup \{\lambda\}$. This holds irrespective of the power of the behavior λ . These results imply that more powerful behavior spaces are consequently less modular. \square

These results show that for memoryless models of intelligence, the more powerful

one tries to make a system, the more difficult it is to avoid destructive cycles. Also, when expanding the behavior repertoire, it is likely that new behaviors will need to be developed along with the old ones, many times at the cost of modularity. These results are of importance to the debate within the AI community on the nature of behavior models. However, in a broader sense, these results also throw up important insight into the role of memory in general architectures of intelligence, which becomes clearer when we compare these results with those from biological systems under conditions of memory deficit, as with the lesions in the parts of a brain or its pathways. This is explored in the next section.

5 Cyclicity in Biological Systems

Stimulus specialization results in a more specific stimulus for a given action. Instead of sitting on all yellow flowers, a butterfly may specialize by going only to marigolds. Response specialization may lead the butterfly to mark flowers it has visited by a characteristic scent, in addition to the original behavior of nectaring. These types of behavior modifications are however not a result of avoiding cyclic conflicts, but of following higher-level goals or motivation. Such actions require some form of memory, which may be a part of a creature's *internal state*. Memoryless systems may not have this type of capability, and those which do not, have actually been observed to exhibit cyclic behavior. This is one of the stronger reasons why pure subsumption robots in the zero-global-memory version may be limited in the complexity of tasks that they can perform. There is a considerable cognitive evidence for spatial memory not only for the scene visible at each moment but

also the broader space that extends beyond the current sensory range [13]. Our results suggest that without such a model for spatial memory, mobile robots will be severely limited in the kinds of the spatial tasks they perform. This explains why autonomous robots that are purely reflexive have been able to achieve little more than obstacle avoidance and wandering over short distances.

Behavior modification of the type discussed in the previous section can be achieved externally or internally. For amnesics who are likely to have no temporal recollection, one could arrange to change the environment, e.g. by putting up a written poster that they can see. In case of the amnesiac Jimmie [29], who addresses the same doctor afresh each time as if he had never met him in his life, one could modify this behavior by keeping around photographs of familiar people. Note the similarity between this and the repeated pick-up behavior of a robot, and how it can be solved by tagging the deposit area as already-deposited. In biology however, absolute absence of memory is never there - even severely amnesic people are capable of having considerable memory, and may be able to learn and make small internal corrections, but they are often themselves unaware of it [15].

Some of the results regarding the loss of power under corrective action in artificial systems may also extend to biological systems. For example, Jimmie using pictures of people requires that the interaction be modified to include photo-taking, matching and adding photos to the list, etc, which reduces the flexibility of the behavior space. Thus, the results that we have formally proved for artificial systems regarding the weakened modularity and higher conflict likelihood in more powerful behavior systems may hold in

biological systems as well, though establishing this formally would require substantially more work.

Another question raised by this study is the nature of memory needed to avoid cyclicity. If a robot transports ten cans and memorizes their locations, then picking them up again can be avoided. However if there are more cans than can be stored in the memory, then some of the can locations may be erased making them liable to be picked up again. With more cans, the likelihood of cycles will increase. Alternately, one may denote some class information, e.g. by tagging the depository room instead so that the robot can ignore all cans in this vicinity. This brings up the need for abstraction in memory, a case that is beyond the scope of this paper.

Learning biological behaviors is a relatively slow process, and stimulus specialization and response generalization mechanisms may not apply as directly to them as they do to the artificial systems. However, instances such as the Pavlovian generalization of salivation by a dog from the sight of food to the sound of the bell involves an association similar to a stimulus generalizing mechanism, where the stimulus for salivation has been generalized to (*presence_of_food* \vee *sound_of_bell*).

One of the differences to keep in mind when comparing artificial memoryless systems such as the purely reactive robots with biological systems such as amnesics or brain-lesioned animals, is that biological systems, even when they do not exhibit memory, are still responsive to their internal motivation. This makes it possible for them to change priorities dynamically, whereas in artificial systems, priorities are pre-assigned by the programmer. The role of motivation in artificial systems is yet to receive the fullest

treatment, and this may hold answers to many of the questions raised here. However, even with a motivation, some memory is clearly needed, without which the system may again exhibit some type of cyclicity.

6 Conclusion

In this work, we have constructed a model for memoryless systems, which can be artificial or biological. Our prime impetus, however, is from the study of the models of artificial behavior, as in the purely reactive robots. The results are also meant for such artificial systems, though the analogical situation in biology adds significance to these, and calls for a possibly independent study of such effects in the biological domain.

Our main result is that in a memoryless system, there is a tradeoff between the likelihood of cyclic conflict and the complexity of the behavior system. Also, these behavior spaces become less modular as they become more powerful, in the sense that if new powerful behaviors are developed in isolation, then cyclic conflicts are more likely to appear when these new behaviors are added to the system. This challenges the modularity assumption, implicitly made by the behavioral model, that intelligence can be developed in independent modules.

Another observation is that the effect of learning may exacerbate the tendency towards cycles, since learning typically produces more general behaviors – e.g. a child climbing up a chair generalizes that behavior to climb up a couch, stair and table etc., thus having generalized the stimulus. Some evidence of this may be found in recent work on robot learning [22].

The other results of this work attempt to offset the negative connotations of the above by identifying a methodology for overcoming such cyclic conflicts. Reliable solutions are unlikely to be based on suppression, where a solution will work only if the stimulus for the conflict disappears by the end of the suppressing behavior. Other methods for avoid cycles involve modifying the behaviors themselves, and this can be done either by specializing the stimulus or generalizing the response of some behavior modules. Unlike learning, which makes the behaviors more powerful, these modifications reduce the usefulness of the behavior spaces. If a robot can pick up a soda can, it should be able to pick up a coffee cup or other similar object. Using stimulus specialization, such a general behavior would be split into many separate behaviors for picking up separate objects. Thus, this solution actually comes at a heavy price.

One possible criticism of the model adopted in this work may be that it uses predicate logic to represent the stimuli and consequences, whereas these may involve some other underlying representation such as fuzzy logic (stimuli may be continuous as opposed to binary), or connections with variable activation levels as in [21]. However, what causes the cycles in these systems is not the representation, but the underlying lack of internal state. For example, if we port the can-pickup example into any of these representations, the conflict does not go away, since the conflict arises at the knowledge level and not at the representation level.

Finally, we note that the cyclicity problems of behaviors also apply to other paradigms in AI such as *planning*. One solution to the cyclicity problems in planning is the use of a meta-level reasoner, an idea which has been used in behavior systems as well [2].

Boden [4] has argued for a greater study of planning as it occurs in animal and human systems. However, behavior models differ from planners in some crucial aspects. Locality of behavior programming makes opportunistic plan-generation automatic, since the relevant behavior is triggered automatically when stimulus becomes sufficiently strong. Also, cycles are much more of a problem in behavior models since unlike the operators in plans, a behavior does not “switch-off-and-die” after execution; if the stimulus reappears, it may re-execute, causing a cycle.

In conclusion, we say that this new approach to the study of artificial memoryless systems can provide important insights into the organization of biological systems with memory, just as the study of amnesics provides important clues into the operation of the brain.

Acknowledgements

We thank Laurie Ihrig, Marvin Minsky, Reid Simmons, Randall Beer, Erann Gat, Harish Karnick, John Bay, Neil Murray, Melinda Gervasio and the anonymous reviewers for their useful comments on this work. We also thank Matt Mason, Avi Kak, B. Chandrasekaran, Suddipto Mukerjee and Henry Hexmoor for encouraging the research that resulted in this paper. We thank Prof. R.E. Passingham and the American Psychological Association (copyright holder) for permitting us to reproduce the figure from [27].

References

- [1] T. L. Anderson and M. Donath, “Animal behavior as a paradigm for developing robot autonomy”, *Robotics and Autonomous Systems*, 6(1 & 2), 1990, pp. 145-168.

- [2] R. C. Arkin, "Modeling neural function at the schema level: Implications and results for robotic control", Chapter 17, *Biological Neural Networks in Invertebrate Neuroethology and Robotics*, Ed. R. Beer *et al*, Academic Press Inc., 1993.
- [3] R. Beer, H. Chiel & L. Sterling, "A biological perspective on autonomous agent design", *Robotics & Autonomous Systems*, 6(1 & 2), 1990, pp. 169-186.
- [4] M. A. Boden, "Artificial intelligence and biological intelligence", *Artificial Intelligence in Psychology*, The MIT Press, 1989.
- [5] V. Braitenberg, *Vehicles: Experiments in synthetic psychology*, The MIT Press, Cambridge, Massachusetts, 1984.
- [6] A. Brodal, *Neurological anatomy*, Oxford Univ. Press, 1976.
- [7] R. A. Brooks, "A robust layered control system for a mobile robot", *IEEE Transactions on Robotics and Automation*, 2(1), 1986, pp. 14-23.
- [8] R. A. Brooks, "Intelligence without representation", *Artificial Intelligence*, 47(1-3), 1991, pp. 139-159.
- [9] J. Connell, *Minimalist mobile robotics, A colony style architecture for an artificial creature*, Academic press Inc., 1990.
- [10] J. Culbertson, *The minds of robots: Sense data, memory images, and behavior in conscious automata*, University of Illinois Press, Illinois, 1963.
- [11] J. C. Eccles, *The understanding of the brain*, McGraw-Hill Book Company, 1977.

- [12] M. P. Georgeff, "Planning", *Annual Review of Computer Science*, 2, 1987, pp. 359-400.
- [13] R. N. Haber and L. Haber, "Why mobile robots need a spatial memory?", *SPIE Vol. 1383 Sensor Fusion III:3-D perception and recognition*, 1990, pp. 411-424.
- [14] M. L. Howe and M. L. Courage, "On resolving the enigma of infantile amnesia", *Psychological Bulletin*, Vol. 113, No. 2, 1993, pp. 305-326.
- [15] L. Jacoby and D. Witherspoon, "Remembering without awareness", *Canadian Journal of Psychology*, 36(2), 1982, pp. 300-324.
- [16] L. Jacoby, J. Toth and A. Yonelinas, "Separating conscious and unconscious influences of memory: Measuring recollection", *Journal of Experimental Psychology: General*, Vol. 122, No. 2, 1993, pp. 139-154.
- [17] D. P. Kimble, "The effects of bilateral hippocampal lesions in rats", *Journal of Comparative and Physiological Psychology*, 56, 1963, pp. 273-283.
- [18] D. Kirsh, "Today the earwig, tomorrow man?", *Artificial Intelligence*, 47(1-3), 1991, pp. 161-184.
- [19] D. Locke, "The indispensability of memory knowledge", *Memory*, The Macmillan Press Ltd., 1971.
- [20] A. R. Luria, *The neuropsychology of memory*, V. H. Winston & sons, 1976.
- [21] P. Maes, "How to do the right thing?" *Connection Science*, 1(3), 1989, pp. 291-323.

- [22] S. Mahadevan and J. Connell, "Automatic programming of behavior-based robots using reinforcement learning", *Artificial Intelligence*, 55, 1992, pp. 311-365.
- [23] A. D. Mali and A. Mukerjee, "Robot behavior conflicts: Can intelligence be modularized?", *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, Seattle, 1994, pp. 1279-1284.
- [24] A. D. Mali and A. Mukerjee, "Metrics for evaluation of behavior-based robotic systems", *Proceedings of the IEEE International Conference on Robotics & Automation (ICRA)*, 1998, pp. 1122-1127.
- [25] P. M. Milner, *Physiological psychology*, Holt, Rinehart & Winston Inc., 1970.
- [26] M. L. Minsky, *The society of mind*, Simon and Schuster, 1986.
- [27] R. E. Passingham, "Memory of monkeys (*Macaca mulatta*) with lesions in pre-frontal cortex", *Behavioral Neuroscience*, 99(1), 1985, pp. 3-21.
- [28] I. P. Pavlov, *Conditioned reflexes*, Oxford university press, Oxford, 1927.
- [29] O. Sacks, "The lost mariner", *The man who mistook his wife for a hat and other clinical tales*, Summit books, 1985.
- [30] A. M. Schneider and B. Tarshis, *An introduction to physiological psychology*, Random House, NY, 1975.
- [31] N. E. Spear, "Behaviors that indicate memory: Levels of expression", *Canadian Journal of Psychology*, 38(2), 1984, pp. 348-367.

- [32] G. A. Talland, *Deranged memory*, Academic Press Inc., 1965.
- [33] R. G. Tharp and R. J. Wetzel, *Behavior modification in the natural environment*, Academic Press Inc., 1969
- [34] N. Tinbergen, "The curious behavior of stickleback", *Scientific American* 187, 1952, pp. 22-26.
- [35] J. Usher and U. Neisser, "Childhood amnesia and the beginnings of memory for four early life events", *Journal of Experimental Psychology: General*, Vol. 122, No. 2, 1993, pp. 155-165.
- [36] A. H. Vera and H. A. Simon, "Situated action: A symbolic interpretation", *Cognitive Science* 17, 1993, pp. 7-48.
- [37] E. K. Warrington and L. Weiskrantz, "The effect of prior learning on subsequent retention in amnesic patients", *Neuropsychologia*, 12, 1974, pp. 419-428.
- [38] S. Zola-Morgan and L. R. Squire, "Medial temporal lesions in monkeys impair memory on a variety of tasks sensitive to human amnesia", *Behavioral Neuroscience*, 99(1), 1985, pp. 22-34.