Probabilistic BDI in a Cognitive Robot Architecture

Darryl N Davis*, Hossein Miri

Department of Computer Science, University of Hull, Kingston-upon-Hull, United Kingdom

*D.N.Davis@hull.ac.uk

***Abstract*- The probabilistic approach to cognition has become an established approach in recent decades. Cognition is better viewed as solving probabilistic, rather than logical, inference problems; i.e. cognition is better understood in terms of probability theory, rather than in terms of logic. This article presents a cognitive architecture used to govern a robot probabilistically. The design and implementation of cognitive architectures is a useful tool for understanding cognition in a situated agent. The Cerno research project extended the CAMAL (Computational Architectures for Motivation, Affect, and Learning) model, by incorporating probabilistic reasoning in its BDI model. Subsequent development of CAMAL has integrated all the valanced affective predicates across the architecture. Extensive experiments with synthetic and real robots demonstrate an improvement in the overall performance, success rate, task effectiveness, and goal achievement of the cognitive architecture.**

Keywords*- *Cognitive Robots; BDI; Probablistic Reasoning; Affect

I. INTRODUCTION

The CAMAL architecture is an example of a general class of integrative cognitive architectures; drawing together a number of threads in Cognitive Science and Artificial Intelligence, such as perception, action, decision making, motivation, affect, and learning. CAMAL (Computational Architectures for Motivation, Affect, and Learning) [1] was developed from ideas incorporated in Guardian [2], ACT-R [3], CRIBB [4], CogAff [5], and the cognitive architectures of Singh and Minsky [6].

CAMAL is, essentially, a UTC (Unified Theory of Cognition) [7] that tries to answer many of the questions that comprise Norman's Cognitive Science agenda [8]. CAMAL uses a slight variation of the reasoning model to be found in a-CRIBB [9, 10]; i.e. a BDI (Belief-Desire-Intention) reasoning model that is driven by an affect model. CAMAL previously has not included any probabilistic reasoning capability, although multi-dimensional metrics (from the affect model) have been used to rank motivations, goals and manage affect [11].

The primary aim of the Cerno research project is to investigate how CAMAL can be extended to reason probabilistically about tasks and domain model objects. In particular, it looked to integrate a probabilistic formalism into its BDI model to coalesce a number of mechanisms. The impetus for this investigation is the considerable evidence that probabilistic thinking and reasoning are linked to cognitive development and play a role in cognitive functions, such as decision making and learning [12, 13, 14]. This leads us to believe that a probabilistic reasoning capability is a vital part of any system that attempts to emulate human intelligence computationally. In other words, probabilistic reasoning is an essential aspect of the process of cognition and, therefore, must be considered in any adequate description of it.

II. COGNITIVE ARCHITECTURES AND ROBOTS

Cognitive robotics (both simulated and embodied) is a growing research field that draws on a number of influences. Mobile robots provide an essential tool when investigating the interaction of cognitive architectures and the physical environment. Robots have been used to investigate many different aspects of artificial intelligence such as mapping and localization techniques [15, 16], robot perception [17] and robot learning [18, 19]. The reported research builds on the previous use of a mobile robot to investigate a specific area of cognitive science known as the anchoring problem [20]. The anchoring problem describes the problem of generating and maintaining links between symbols and perceptual data.

Hybrid architectures seek to avoid the disadvantages of their component architectures, whilst retaining all their benefits. A commonly used hybrid architecture is the reactive-deliberative architecture [21]. Here (see Figure 1) we present an architecture consisting of several different elements including a set of low-level robot actions; a reactive component built from many different reactive behaviours; a belief-desire-intention (BDI) schema; a distributed model of affect; an association construct; a domain model; and a motivational blackboard that links all these subsystems together.

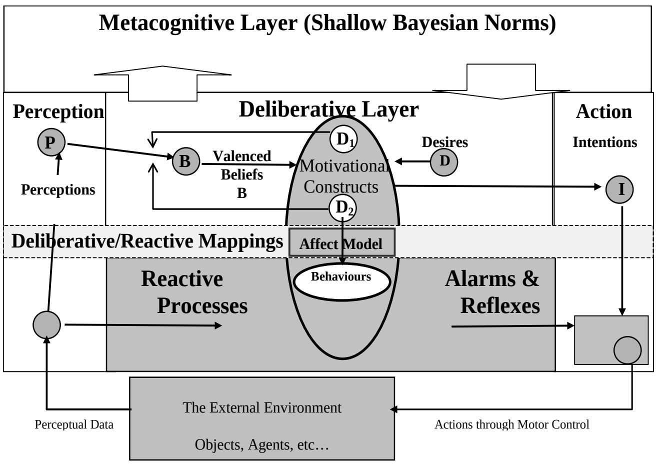

Fig. 1 Schematic of the Cerno Architecture

Most pertinent to the overall aims of the current research, is the mobile robotics work of Stoytchev and Arkin [22].

DOI: 10.5963/IJCSAI0203001

Their architecture combines three components: deliberative planning; reactive control; and motivational drives. They identified that the mapping of high-level deliberative commands onto a reactive controller solved some of the problems associated with purely reactive control. The approach taken in Cerno (plus robo-CAMAL and other CAMAL variants) offers a solution, at least in the limited experiments performed to date. The BDI schema used with the association mappings offers a generic solution that can be tailored to specific environments and tasks, through adaptation and learning of the domain model. The motivational model used in CAMAL is developed from a goal based model of emotion and operates at both the deliberative and reactive level.

III. CAMAL AND CERNO-CAMAL

The belief-desire-intention (BDI) model [23] is a schema that calculates the actions of an agent based on its beliefs and its desires. A belief is a statement about the confidence of a proposition. The confidence the agent can have in a belief can vary. In the BDI model, beliefs are based on input from the agent's perceptual system, and its previously held beliefs. The agent's desires are a set of goals, which the agent wishes to achieve. The agent's current desires are based on its internal state, possibly its emotional state, and its previously held desires. Through a coupling of the agent's beliefs and its desires, CAMAL generates a set of intentions or plans to achieve its goals. For example if the agent has a goal to hit a ball, and the perceptual system generates the belief that there is a ball to the right, the agent can then implement a set of plans to turn the agent right and move forward.

The CRIBB (Children's Reasoning about Intentions, Beliefs and Behaviour) model was developed to investigate reasoning in young children [24]. This schema was implemented as a computer model to simulate knowledge and the inference processes of a child solving problems [4]. Earlier research saw a variety of CRIBB using affect (a-CRIBB) developed [9, 10, 11]. For robo-CAMAL, a-CRIBB was developed with a major difference to the standard BDI model in that different qualitative degrees of belief are ascribed to belief statements, according to the source of the belief. A preference operator allows discrimination between beliefs that are based on assumption, perception and deduction. Perception can be further sub-divided to include direct perception and indirect perception (i.e. from another agent that has some degree of trust associated with it). Hence beliefs can be ordered according to the degree of trust in them. The current version of CAMAL (and hence the Cerno variant) develops this further with the use of (quantitative) belief metrics.

CAMAL makes use of a hybrid reactive-deliberative architecture (Figure 1) based on the control system approach to mind [25]. A control state is a behavior internal to an agent. Control states can exhibit external behaviors (such as obstacle avoidance) or, reflect and control internal states (such as beliefs). In essence, control states can be a number of things, such as beliefs, desires, motivators, etc. They can be considered as general behaviors within an agent, and therefore can be modified or updated at any time by some other processes within the agent.

Scientific evidence highlights the fundamental role of affect in rational and intelligent behaviour [26]. Since there appears to be a requirement for something analogous to affect in artificial cognitive systems, it has been conjectured to use affective control states which makes affect the basis of a consistent control language across a cognitive architecture [9, 11, 25, 27, 28]. Further research proposes that motivation also plays a fundamental role in a variety of cognitive functions [29]. From the viewpoint of mind as a control system, motivation can thus be thought of as a control state. Furthermore, motivators are dispositions and tendencies to assess situations and respond to those situations and assessments in a certain way. They can provide a context and impetus for reasoning about events, and also a basis for goal-directed behaviors. They are, therefore, often used as a generic framework that draw all these control states (beliefs, goals, etc.) together.

It now seems reasonable to conclude that motivation and affect can be co-joined in perception and cognition. Motivations are essentially more encompassing than goals since they include not only the descriptions of goals but also an affective context for those goals. Furthermore motivators are often used as a generic framework since they can draw all the control states of a cognitive architecture together. Hence, the use of control states to develop cognitive architectures leads to the use of affective and motivational control states. The following sections briefly introduce some of the main components of the CAMAL cognitive architecture.

A. *Probabilistic -BDI Model*

One of the limitations of the earlier BDI model was the lack of any explicit mechanism to express degrees of belief. CAMAL represents beliefs as categorical states; prioritized by a CAMAL preference model, using belief preference operators. Given that our current work on this theme presents an affect- and affordance-based core for mind, it seems reasonable to conjecture that beliefs, too, should be grounded in the use of affect. The aim is for this mechanism to be consistent across different domains, tasks, and levels of processing. This will be compatible with the way in which the affect and motivational models operate throughout CAMAL, having an associated affective magnitude that can fluctuate according to success or failure associated with that element. Effectively, affect will serve as a decision metric and affective values as a currency across the entire architecture. We, therefore, represent belief statements as graded states, with the Extended Belief Structure (EBS) using clauses of the form:

B. *belief (Descriptor, Source, Time, Degree Belief)*

The numerical reasoning algorithms underpinning the EBS associate a probability value *DegreeBelief* with every belief statement in CAMAL. This defines the degree to which the belief statement is believed to be true. This enables the computation of degrees of belief, and using the BDI and affect models to determine the agent's intentions, actions, and behaviors. It will also allow the entire BDI

DOI: 10.5963/IJCSAI0203001

model to run using numeric affective values to prioritize choices made over a current belief set.

CAMAL's Probabilistic-BDI model provides a method to control the flow of information through the deliberative component. It determines the intentions, actions, and behaviors of the cognitive agent based on its beliefs and desires. For example, the Domain model (described later in this section) contains belief statements about the assumed reliability of belief sources (perception sensors, deduction or assumptions):

degree_of_belief (perception, 0.9).

degree_of_belief (deduction, 0.75).

degree_of_belief (assumption, 0.5).

The belief handling mechanisms (in the Deliberative Layer of Figure 1) can therefore use (or modify) these predicates when reasoning over beliefs to generate statements such as the following statements taken from an instance of CAMAL running in a test-bed (*robotworld*) populated by other robots and a ball (and used in the third experiment in section IV):

belief (environment (sparse), assumption, 1, 0.5);

belief (cycles (2), deduction, 2, 0.75);

belief (energy (100), perception, 2, 0.9);

belief (found (robot1), perception, 2, 0.9);

belief (near (ball), perception, 2, 0.9).

C. Domain Models

Underpinning the Probabilistic-BDI, affect, and motivational models, CAMAL makes use of domain models. The cognitive architecture shown in Figure 1 (Cerno) contains three processing layers: reactive, deliberative, and metacognition. The latter layer is more extensive in the full CAMAL implementation. It is vital that relevant information about its attributes and properties along with its surrounding environment be incorporated into its cognitive architecture. This incorporation is achieved by the use of its (propositional) a-priori domain model.

This model (which the user supplies as a set of propositions) defines how CAMAL should reason about its intended test-bed. CAMAL is meant to be a generic architecture that can be initialised to reason about a specific set of tasks in a specified environment. Using a domain model is a method of instantiating the relevant information about a cognitive agent's attributes and its surrounding environment into the cognitive architecture. The domain model, once loaded, initialises the motivational blackboard, parameterises the metacognitive layer, and is then distributed across the cognitive architecture, at both the reactive and deliberative levels. It provides information on the agent's attributes and its surrounding environment. It defines the types of objects to be found within the agent's environment (physical or simulation). It defines some of the possible beliefs that are most relevant to the agent. It defines the relationships between the stated beliefs. It defines the goals that the agent can have. It provides a list of all the possible actions the agent can undertake; and it provides the objects' perceptual profiles (i.e. the information required to recognize an object), etc.

Statements about the agent's environment at the deliberative level pertain to the possible beliefs the agent can hold. These constitute the beliefs used in the BDI model. These statements are presented on the motivational blackboard at the deliberative level in Figure 1, unlike the environment domain model which had components present at both the reactive and deliberative levels. These statements are placed on the motivational blackboard, and provide information on the goals that the agent can have and the actions that the agent can take. These constitute the desires and intentions used in the BDI model.

Goals take the following form:

Goal (Descriptor, Success Condition, Importance).

The Descriptor is a description of the agent's goal. The Success Condition is the belief descriptor required for the goal to be achieved. The importance is an affect value detailing the goal's affordance, meaning how important the goal is to the agent.

CAMAL, once running, is allowed to modify the loaded domain model, and so tune the architecture to the specific environment it finds itself in. Example goals are given in the explanation of the association structure in the following section.

D. Affect & Motivational Models

Affect is defined in terms of information processes and representational structures across the cognitive architecture. It is qualitatively defined as negative, neutral, or positive, and can be mapped numerically (as a *valance*) over the interval *[-1.0, +1.0]*. The use of affect and affective control states makes affect the basis of a consistent control language across the architecture. It allows external events and objects to take valanced affordances, and internal mechanisms to be prioritized via valanced processes. Adding affect results in more effective processing and task management [9, 11].

The affect model distributes affect values across the entire cognitive architecture, rather than having a centralized emotion module. This is made possible by the use of associations. An association is a construct that contains a complete BDI combination, as well as an associated affect value (*insistence*). This association value is one of the main factors in deciding (at the deliberative level) the next action of the cognitive agent. Associations take the following form:

Association (Belief, Desire, Intention, Insistence).

Example associations from CAMAL running in the *robotworld* domain, introduced above, include:

Association (environment (sparse), find (ball),
Reactive (true, true, method5), 0.95).

Association (near (ball), hit (ball),
Reactive (false, true, method_hit), 0.95).

DOI: 10.5963/IJCSAI0203001

association (found (robot1), approach(robot1),

reactive(false, true, method_approach), 0.95).

association (found (robot1), track(robot1),

reactive (true, true, method4), 0.05).

association (found(robot1),track(robot1),

reactive (true, true, method5), 0.835).

The architecture maintains various beliefs about the environment it is operating in, and several possible goals that relate to that environment. It also has a number of different plans of actions (or intentions or behaviors) that correspond to specific reactive sub-architectures that can be activated (as in the last two example associations given above). Associations effectively provide a method for a cognitive agent to keep track of all these possible belief-goal-action combinations, as well as containing a key value that indicates their relevance and significance to the agent at a given time—their association value. These combinations detail the correct action required to achieve a specific goal given a specific belief.

Given the example belief and association predicates above, all associations have their belief basis substantiated. CAMAL would reason that the first association has already succeeded as the *find (ball)* goal has already succeeded (belief *near (ball)* exists). If the chosen goal were *track (robot1)* then the second of the two relevant associations would be chosen as it has a higher *Insistence* value. However, as the following text describes, CAMAL chooses a motivation that is a combination of affordances from current belief, goal and association statements.

As part of the distributed affect model, the motivational model contains multiple affective values in the structures of its motivational blackboard. The motivational blackboard used in CAMAL is akin to the structures found in Blackboard Systems [30], and acts as a global workspace. It is a global structure that holds all the relevant information about an agent's environment, attributes, properties, beliefs, desires, current state, previous states, etc.. This blackboard, analogous to Baars' Global Workspace [31] is potentially accessible by all the processes of an agent. The knowledge and information held on the motivational blackboard can be divided into several distinct areas: beliefs that the agent can have about its environment; beliefs that the agent can have about its own internal state; desires that the agent can have about the objects in its environment; intentions that the agent can have to achieve its goals; associations that are used to manage the BDI and affect models; and a motivator that contains the result of the operation of knowledge sources. The reasoning focus in CAMAL is often on the motivator as a representational form that enables perception, affect, cognition, and behavior to come together and interact. In effect, a motivator is a representational schema that is used as a generic framework to bring together aspects of cognitive processing, such as perception, affect, and behavior. Motivators are used to manage and manipulate the affective and motivational control states presented at the 'deliberative' level. They may also trigger appropriate intentions, actions, and behaviors at the 'reactive' level, such as selecting a suitable reactive sub-architecture. Motivators can, thus, be thought of the result of the operation and execution of the various (reactive, deliberative and meta-cognitive) knowledge sources on a motivational blackboard. Motivators take the following form:

Motivator (Goal, Association, Time, Determinism, Cycle, Intensity)

The *Intensity* element is an affect value that defines the importance of the motivator to the agent. The schema also contains the agent's chosen goal, plus the appropriate association selected to achieve that goal, which obviously contains the intention of the agent as well (as highlighted above). The *Cycle* element gives the number of cycles the reactive component should run for. The *Time* element defines the deliberative cycle when the motivator was first generated. The *Determinism* element is Boolean. If set to true, the reactive component would override any reactive/ behavioral failure in achieving a goal or other conditions until the end of the given number of cycles, or until the goal is met. If set to false, the reactive component would return the first possible action (including behaviour failure) that the deliberative-reactive interface selects, based on the appropriate goal-association combination.

Once a goal and a relevant association are chosen, the motivator-update knowledge source of the motivational blackboard uses them to update the motivator. For the running *robot world* domain example, given the set of beliefs and associations listed above, with all goals defined in the associations also present on the motivational blackboard, then CAMAL would generate a motivator of the form:

Motivator (goal (approach (robot1), near (robot1), 0.68),

Association (found (robot1), approach (robot1),

Reactive (false, true, method approach), 0.95),

2, false, 10, 0.725).

This means that CAMAL will attempt to satisfy the goal *approach(robot1)* using the reactive level (see Figure 1) behavior *method approach* with sonar switched off and vision on (as defined by the atoms *false* and *true* in the reactive behaviour definition in the association). The reactive behaviour will run for 10 (reactive level) cycles but can be interrupted if the reactive behaviour fails during any of those cycles (as defined by the atom *false* between the *Time* and *Cycles* fields). The motivator was created in deliberative Cycle 2 with an *Intensity* as given in the last field. It should be noted that the architecture shown in Figure 1 (as with all CAMAL variants) is asynchronously parallel. The overall architecture has different run times and cycles at the meta-cognitive, deliberative and reactive levels. Hence the explicit separation of deliberative and reactive time and cycles in the deliberative level motivator, and the use of separate predicates to refer to each at the deliberative level..

DOI: 10.5963/IJCSAI0203001

E. Operational Overview

Cerno is variant of CAMAL that uses reactive behaviours associated with robo-CAMAL [20] and the first CAMAL architecture to incorporate the Probabilistic-BDI schema. It does not include the meta-deliberative reasoning of CAMAL [32] but simply a set of statements that define how it should reason. These *norms* (given in the Domain Model) define how the algorithms underpinning the combination of affective valances in BDI model are combined. Unlike CAMAL it cannot reason about these or modify how earlier failed or unused motivators are revisited when current motivators start to fail.

The inclusion of degree-of-belief in the structure of its belief predicates enables the architecture to select a focused belief set that reflects its current activities, as highlighted by actions, objects, and agents referenced in a current motivator. The motivator enables goal revision and the selection of the next goal based on goal importance, current beliefs and goal history. The deliberative processing of these constructs allows the selection of an appropriate action related to specific objects and tasks. This, in turn, drives motivator revision using the association construct, which in turn enables belief-desire-intention combinations to be ranked based on the likelihood of their success (association insistence values). The goal importance, association insistence, motivator intensity, and degree-of-belief are all underpinned by affordances; i.e. they are all consistently grounded in affect. Together, they allow motivators to persist or be updated by new goals, associations, etc.

At the reactive level, perceptual data from the agent's sensors in the test-bed are passed to the deliberative layer. This perceptual message is parsed and posted to the motivational blackboard. Whether in simulated or robotic test beds, sensory information is mapped onto belief structures. Belief affordances (degrees of belief) define the degree to which the belief statement is believed to be true. The belief-update module uses the new information to modify its belief set. The motivator assessment modules then uses the updated belief set to determine the success of the current motivator if the current goal has been achieved. This then leads to goal updating and the selection of a new goal set. The association-update then uses the new belief and goal sets to determine the relevant action or intention. The affective insistence measure allows the control of external behavior through the building of associations that link beliefs, goals, and intentions.

IV. EXPERIMENTAL RESULTS & OUTCOMES

A succession of experiments was carried out to evaluate Cerno's overall performance, success rate, task effectiveness, and goal achievement.

A. Goal Achievement in a Virtual World

A succession of experiments was carried out to evaluate Cerno's overall performance, in terms of goal success and failure, and consequently task effectiveness.

These experiments used a graphical predator-prey world created using SWI-Prolog. This tile world is populated by objects and operators beyond the control of Cerno that affect the world. These objects include static opaque obstacles, several edible spheres (blue energy objects), red prey agents, green predator agents, and a (white) Cerno agent. Five experiment scenarios, of increasing complexity, were used in order to investigate probabilistic reasoning and behaviour selection. The simplest investigated collision avoidance and navigation; where Cerno avoids every object, whether an obstacle, energy sphere or other agent. Scenario Two involved a self-maintenance task where Cerno maintains its energy level through the finding and consumption of energy spheres. Scenario Three involves task-oriented goals where Cerno finds and herds prey agents. Scenario Four involves task-oriented goals where Cerno finds and attacks predator agents. Scenario Five builds on the first four involving both self-maintenance and task oriented goals. In this final scenario, Cerno acts as a virtual sheepdog feeding off energy spheres, herding prey agents, and protecting prey against predator agents.

The objective was to assess the efficacy of the Cerno architecture over the original CAMAL in a tangible manner, by using success and failure counts. Efficacy is measured in quantitative terms here as greater number of successes and lower number of failures. These tests are performed to ensure that the use of the EBS, and assumed degrees of belief, infer posterior probabilities correctly, assign them to the appropriate belief descriptors, and reason probabilistically about the number of objects and their instances that may be present in the environment. Cerno was allowed to operate in the described virtual world with a varying number of energy objects, prey agents, and predator agents. Each experiment was run for 10 deliberative cycles. This was repeated with and without the presence of the opaque environmental obstacles.

In the predator-prey tile world, the term success refers to one of the following.

1.The set goal (defined in the domain model) is achieved when the new beliefs explicitly state it. This situation is known as *Explicit Goal Success*.

2.The set goal (defined in the domain model) is achieved when new beliefs state that the negation of the goal negation has been achieved. CAMAL allows the negation of Beliefs in Goal descriptor. The negation of this negated Belief clause (either as *not (not (Belief))* or more simply *Belief*) in the subsequent Belief set is known as *Double Negation Success*.

3.The intention (defined in the association construct) is accomplished, when the new beliefs state that. This situation is known as *Explicit Intention Success*.

4.The avoid-collisions intention (defined in the association construct) is accomplished when the new beliefs state that. This situation is known as *Avoid-Collisions Success*.

These results (Figure 2) serve to demonstrate a clear increase in the number of successes. This directly supports the overall advantage of the Cerno architecture over the earlier version of CAMAL in terms of success counts.

DOI: 10.5963/IJCSAI0203001

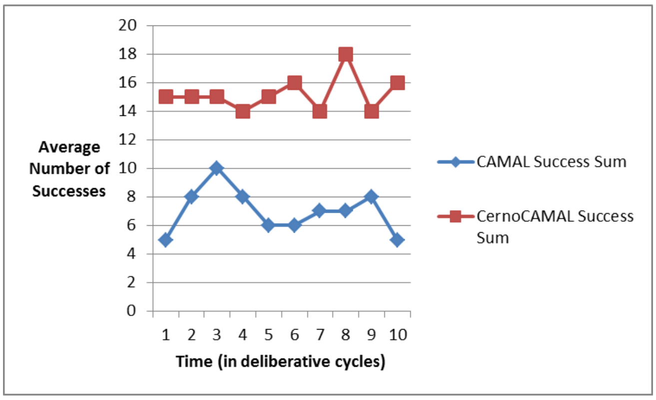

Fig. 2 Overall increase in the number of successes

Similarly, the term failure refers to one of the following:

1.The set goal (defined in the domain model) is not achieved, and explicitly stated in the new beliefs. This situation is known as *Explicit Goal Failure*.

2.The set goal (defined in the domain model) is not achieved, with new beliefs that state that the negation of the goal has been achieved. This situation is known as *Goal Negation Failure*.

3.The set goal (defined in the domain model) is achieved, but on the wrong object as stated by new beliefs. This situation is known as *Wrong-Object Goal Failure*.

4.The intention (defined in the association construct) is not accomplished, explicitly stated as new beliefs. This situation is known as *Explicit Intention Failure*.

5.The avoid-collisions intention (defined in the association construct) is not accomplished, explicitly stated in the new beliefs. This situation is known as *Avoid-Collisions Failure*.

These results (Figure 3) serve to demonstrate a clear decrease in the number of failures. This directly supports the overall advantage of the Cerno architecture over the earlier version of CAMAL in terms of failure counts.

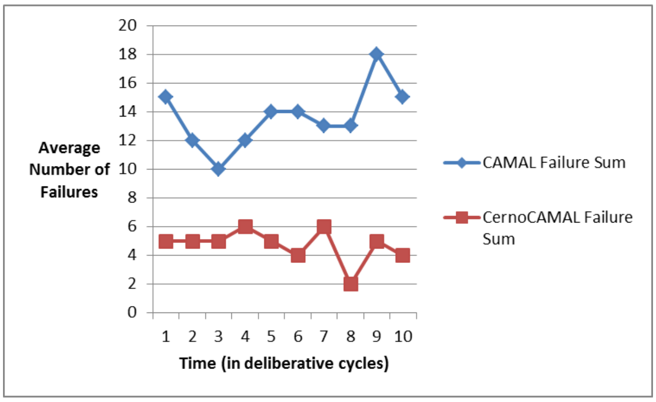

Fig. 3 Overall reduction in the number of failures

Summarily, the term task effectiveness describes if overall the tasks (goals and intentions collectively) were successfully completed. This is inferred based on the increase in the number of success counts and the reduction in the number of failure counts. The two graphs clearly show a decrease in the number of failures along with an increase in the number of successes. This outcome indicates the overall advantage of the Cerno architecture over the earlier (non-probabilistic) CAMAL in terms of success / failure counts, and also task effectiveness

B. CAMAL and Cerno Experiments in ARIA MobileSim

This section presents the ARIA [33] robotic experiments carried out using the Cerno cognitive agent in the MobileSim test-bed. A number of experiments were performed over a number of cycles, and relevant internal variables and statistics were recorded. From the obtained results, the two cognitive architectures of (pre-Cerno) CAMAL and Cerno can be compared and contrasted.

These experiments investigate Cerno's ability to adapt in the dynamic and uncertain MobileSim environment. It is dynamic as it is inhabited by moving robots besides the P3DX that is running the Cerno cognitive architecture. It is uncertain as there is built-in sensor noise and added random error in the simulation test-bed. To investigate adaptability some experiments were paused and modifications made to the number of specific objects or robots were made. Upon resuming experimentation, the findings about the probable number of objects and their instances should, again, be close to the actual numbers set by the experimenter during the pause.

A succession of experiments was carried out to evaluate Cerno's overall performance in terms of task effectiveness, measured by means of goal success and failure. The objective here was to assess the efficacy of the Cerno architecture over the original CAMAL in a tangible manner, by using success and failure counts. In other words, efficacy is measured in quantitative terms here, as the greater number of successes and lower number of failures. A general sample set of obtained experimental results are presented in Figures 4 and 5 to enable the comparison of the two cognitive architectures.

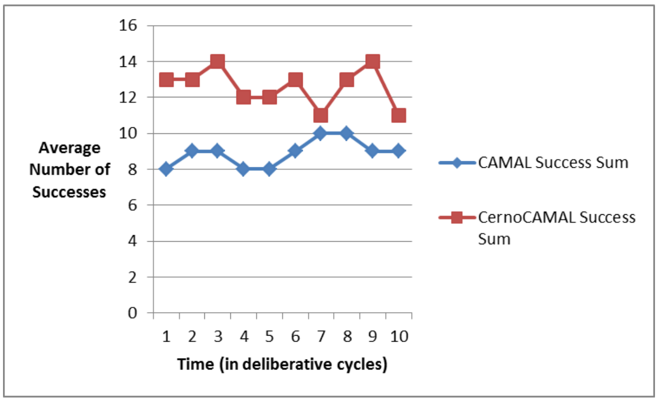

Fig. 4 Overall increase in the number of successes

In the ARIA MobileSim world, the term success and failure are defined as given for the Virtual world experiment in Section IV A above.

These results (Figure 5) serve to demonstrate a clear decrease in the number of failures that occurred, directly supporting the overall advantage of the Cerno architecture over CAMAL in terms of failure counts.

DOI: 10.5963/IJCSAI0203001

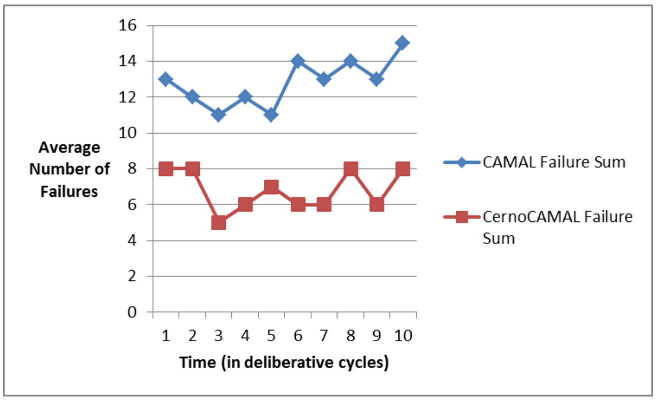

Fig. 5 Overall reduction in the number of failures

Summarily, the term task effectiveness describes if overall the tasks (goals and intentions collectively) were successfully completed. This is inferred based on the increase in the number of successes and the reduction in the number of failures. The two graphs clearly show a decrease in the number of failures along with an increase in the number of successes that occurred. This outcome demonstrates the overall advantage of the Cerno architecture over (the pre-Cerno) CAMAL in terms of success and failure counts and also task effectiveness.

C. Robo-CAMAL and Cerno Experiments in ARIA MobileSim

The Robo-CAMAL research project [20] specifically investigated the anchoring problem in a mobile robot running a simplified version of the CAMAL cognitive architecture. The anchoring problem addresses the linking of perceptual data about objects and events to symbolic representations of those objects and events. In other words, anchoring is the establishment and maintenance of a correspondence from sensory data to propositions denoting (domain model) objects identified from within the sensory data, and actions upon those objects.

Due to the different research goals and objectives, experiments performed with Robo-CAMAL are not directly mapped onto Cerno. Only two significant experiments carried out with Robo-CAMAL have been identified to ascertain whether an improvement or rectification has been accomplished in the process of integrating probabilistic reasoning ability into CAMAL. The rationale behind this selection was the fact that the first experiment highlighted a shortcoming in Robo-CAMAL. It is therefore, a sensible point of reference and comparison between the two architectures. The second experiment highlighted a strength in Robo-CAMAL and is, therefore, an appropriate point to ensure that Robo-CAMAL was not compromised by the inclusion of belief affordances. The obtained experimental results are summarized into two graphs (Figures 6 and 7) that clearly illustrate the argued points.

These experiments used an Amigobot [33] robot; this has two driven wheels with a rear stabilising wheel. It senses the environment through an array of eight sonar sensors; four facing forward; one placed on either side; and two at the rear. In addition to this standard Amigobot configuration, an omnidirectional vision system has been attached. This allows the robot a 360 field of vision (see [20]).

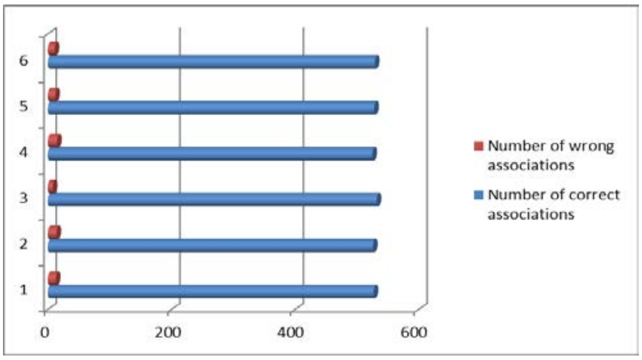

Fig. 6 Negligible number of wrong associations for faulty sensor experiment

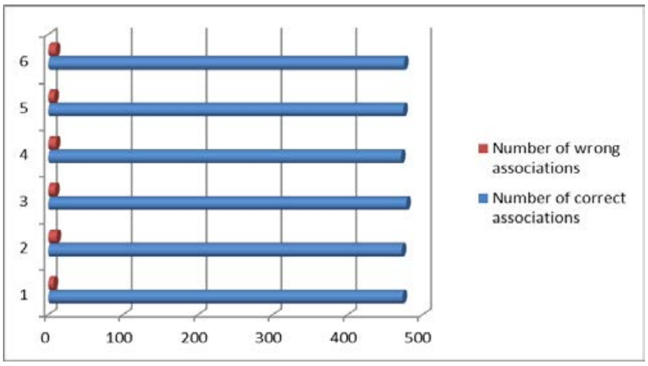

Fig. 7 Negligible number of incorrect associations for adaption experiment

Cerno was used to control the robot in six variations of the same environment for five minutes each. The experiment was repeated three times for each environment. Table 1 shows the six possible environment combinations used for the experiment.

TABLE I ENVIRONMENT COMBINATIONS FOR AMIGOBOT EXPERIMENT

Environment	Object(s)
1	Blue Ball
2	Black Robot
3	Red Robot
4	Blue Ball + Black Robot
5	Blue Ball + Red Robot
6	Blue Ball + Black Robot + Red Robot

Each experiment produces a number of associations based on beliefs derived from the sonar and vision percepts. Robo-CAMAL generated a significant number of associations concerning the black robot when it was not actually present. This result highlighted the difficulty associated with using real sensors. Robo-CAMAL constructed the wrong belief *found (black_robot)* as the vision system had incorrectly identified the object as a black robot when in fact a blue ball was present. It was clear that the cause of this failure was the primitive vision system incorrectly identifying a blue ball as a black robot. While we are developing more sophisticated vision systems, this failure is symptomatic of the type of faulty sensor problems that real autonomous robots need to manage. Hence we ran

DOI: 10.5963/IJCSAI0203001

Cerno with its improved belief mechanisms to see how if it faired.

With Cerno controlling the robot, there were a number of (wrong) associations indicated that the two domain model objects were mistaken, but in the large context of the experiment, this number was negligible. The summary results are shown in the graph of Figure 6. The different environments are marked using numbers one to six. The minuscule number of wrong associations (2.3%) are plotted against the total number of generated associations, to show that the percentage of failures was negligible. %). The earlier Robo-CAMAL in similar experiment averaged 16%, so the probabilistic belief predicate and reasoning change have shown a good improvement.

The other significant successful experiment performed using Robo-CAMAL was an adaption experiment. An important point to note here is that in the context of Robo-CAMAL, adaptation referred to its ability to modify goal importance, and therefore selection, to reflect changes in its environment. The obtained experimental results showed that Robo-CAMAL had the ability to adapt to a variable environment, and attempted the goals it believed achievable at the right time. Similar to the previous set of experiments, it would be instructive to reflect whether the changes in the Cerno architecture might have compromised this capability of Robo-CAMAL.

For this experiment, Cerno architecture was instantiated with three goals:

Hit (blue_ball) & hit (red_robot) & hit (black_robot).

Plus the correct associations were given to the architecture at start up, to determine whether it can modify its goals to reflect changes in its environment. Cerno was then allowed to run three minutes in a variable environment. The environment contained the six possible combinations used for the previous experiment, listed in Table 1. These combinations were changed at intervals of one minute.

Each experiment produced a number of associations. In Robo-CAMAL, most of the generated associations reflected the changes made at one-minute intervals. The summary results are shown in the graph of Figure 7. In Cerno this, too, succeeded based on the huge percentages of correct-to-incorrect associations that showed Cerno had modified its desires to reflect the changes made to its environment. There were a number of (wrong) associations that indicated adaptation took at times up to a whole minute, but in the large context of the experiment, this number was negligible (1.95%). The earlier Robo-CAMAL in similar experiment averaged 22% errors, so the probabilistic belief predicate and reasoning change have shown a good improvement.

V. DISCUSSION

Cerno as a Probabilistic-BDI cognitive architecture has pursued a perspective informed by affective and motivational control states, rationalized by cognitive models of probabilistic reasoning. It presents a vigorous affect- and affordance-based core for mind, as the Probabilistic-BDI model is now valanced via affective values and affordances. This allows the entire BDI model to run using numeric affective values to prioritize choices over the current belief set. The current CAMAL research has now taken a number of new directions, as researchers pursue their own agenda. These new directions take the original design associated with the overarching CAMAL architecture, together with the concept of an underlying affect and affordance mechanism that can be used to compare process priority or rank goals and intentions but reframe the research according to specific interests or needs [11]. Using a larger variety of mobile robot types, along with a new vision system will be the next steps, resulting in a deeper perceptual anchoring model. This new perceptual anchoring model could combine the improved sensors of the new robot with neural learning mechanisms and may, therefore, address some of the issues raised by Barsalou [34].

There clearly existed a need to address and incorporate probabilistic reasoning and inference in CAMAL. The primary aim of the Cerno research project was to tackle this need with the formal probability and Bayesian theories. This research, therefore, attempted to address the following specific research questions in the current cognitive architecture under investigation:

- Can Cerno reason probabilistically by exploiting the proposed EBS? Can the integration of the proposed EBS facilitate probabilistic reasoning and inference in Cerno?

In light of the overall collection of experimental results, Cerno can reason probabilistically by exploiting the proposed EBS. In other words, the integration of the described EBS facilitated probabilistic reasoning and inference in Cerno. This was specifically validated and confirmed, as correct operation of Cerno in terms of probabilistic object/instance reasoning was tested comprehensively.

- Can the Probabilistic-BDI model run compatibly with the affect and motivational models, and affective and motivational valances used throughout the whole architecture? Can this ensure a consistent metric across all aspects of affect, reasoning, and domain model management?

The correct results and expected operations of the processes indicated the compatibility of the Probabilistic BDI model with the affect and motivational valances used throughout the architecture. This provides a consistent control language for ordering propositions, selecting goals, constructing a plan of action, forming a focused belief with an updated degree of belief, and prioritising processes. It ensures a consistent metric across all aspects of affect, reasoning, and domain model management.

- Can the probabilistic deliberation results be used for computing changing degrees of belief given apriori, and subsequently using the Probabilistic-BDI, affect, and motivational models to determine the agent’s intentions, actions, or behaviours?

The performed tests confirmed and validated that the Cerno’s probabilistic reasoner can deliberate using the EBS

DOI: 10.5963/IJCSAI0203001

and assumed degrees of belief, infer posterior probabilities correctly, assign them to the appropriate belief descriptors, and reason probabilistically about the number of objects and their instances that may be present in the environment.

- Can the Cerno cognitive architecture be applied to virtual and physical cognitive agents using synthetic test-beds and mobile robots?

The succession of experiments in simulation and robotic test-beds, by successfully applying the Cerno cognitive architecture to virtual and physical cognitive agents, showed improvements and increased efficacy in Cerno's overall cognitive performance, as well as specific achievements in light of the overall collection of experiments.

The current CAMAL research has now taken a number of new directions, building on the results from associated researchers pursue their own agenda. These new directions take the original design associated with the overarching CAMAL architecture, together with the concept of an underlying affect and affordance mechanism that can be used to compare process priority and rank goals and intentions but re-frame the research according to specific interests or needs [11]. The results and conclusions from these CAMAL related projects feed back into the theory underpinning CAMAL and the resulting CAMAL implementations.

Using a more sophisticated mobile robot such as a P3DX [33] along with a new vision system and camera could be the next step, resulting in a more accurate object identification and consequently a deeper perceptual anchoring model. This new perceptual model could combine the input from the improved sensors of the new robot with a-priori information included in the domain model. The new robot would be more adaptable and capable of working in unknown and uncertain environments. This is loosely related to the UK Computing Research Committee's Grand Challenge Number Five: Architecture of Brain and Mind [35]. GC5 is a multidisciplinary attempt to understand and integrate natural intelligence and high-level cognitive processes at various levels of abstraction. The aim is to demonstrate the results of our improved understanding in a succession of increasingly sophisticated working robots.

The next step could be regarding the manual incorporation of shallow Bayesian metacognitive norms. Currently, Cerno does not deliberate to determine which norm should be used in the motivational blackboard. This means that the Bayesian norms have to be pre-programmed prior to start-up (hand-coded and defined in the domain model). This manual incorporation could be improved upon by constructing more norms and deliberating to choose one that has already yielded greater task effectiveness in the past.

In addition to the above two specific ways of improving the Cerno architecture, there are still plenty of open issues in cognitive architectures research that deserve attention and effort from researchers in the area, despite the many advances that have occurred during almost four decades of research and work. An outstanding issue is that each existing cognitive architecture exhibits many of the capacities described in this thesis and elsewhere, but few support all of them. However, a cognitive architecture as a UTC was defined as a single set of mechanisms and processes for all cognitive behaviour. The research community should perhaps devote more resources to trying to unify the existing capacities and capabilities into one universal and comprehensive framework of cognition.

Moreover, methods for the evaluation and assessment of cognitive architectures and their cognitive abilities could be broadened to include more realistic terrains. Metrics like those used here for experimentation purposes are necessary, but not sufficient to provide an accurate way of comparing and contrasting competing architectures and cognitive systems. Despite evaluating various cognitive performances in different test-beds, more complex environments must be created, both physical and simulated, which exercise these cognitive capabilities and provide realistic opportunities for measurement [36]. Experimental comparisons among competing architectures can play an important role in measuring key variables in unbiased and informative ways. On the positive side, we now have over four decades worth of experience and development with constructing and using a variety of cognitive architectures for a wide range of problems and terrains

VI. CONCLUSIONS

Cognition is better viewed as solving probabilistic, rather than logical, inference problems; meaning cognition is better understood in terms of probability theory, rather than in terms of logic [37, 38]. The probabilistic approach to cognition has, therefore, become an established approach in recent decades –something that this body of work took advantage of.

This paper presents a Probabilistic-BDI affective-motivational cognitive architecture. There are many views on what constitutes a cognitive architecture or the place of affect and motivation in a cognitive architecture. Cerno pursued a perspective informed by affective and motivational control states, rationalized by a cognitive model of probabilistic reasoning using degrees of belief. Any thesis that deals with cognition and cognitive architectures needs some explanation as to its scope and focus. The scope and focus of the current cognitive architecture under investigation were to extend the original overarching cognitive architecture of CAMAL to enable it to reason probabilistically about domain model objects through perception. We have integrated probabilistic formalism into the BDI model to coalesce a number of mechanisms, in line with the Gibsonian affect and affordance theory, as well as Davis's theory of affect [10, 11, 39].

The succession of experiments in simulation and robotic test-beds established improvements in cognitive performance through the adoption of probabilistic inference. In applying the Cerno cognitive architecture to the Robo-CAMAL platform similarly saw a significant improvement, with the number of wrong associations dramatically reduced. The latest version of CAMAL effectively presents a vigorous affect and affordance based core for mind. The

DOI: 10.5963/IJCSAI0203001

Probabilistic-BDI model is now valanced via affective values and affordances, allowing the entire BDI schema to run using numeric affective values to prioritize choices over the current belief set. Since affect is used across the cognitive architecture as a decision metric, affective values can be thought of as a currency. The BDI model that lacked an affective decision metric consistent with the affordances used in the affect and motivational models, is now grounded consistently in the use of affect.

REFERENCES

[1] Davis, D. N., Control States and Complete Agent Architectures. Computational Intelligence, 17 (4): 621-650, 2001.

[2] Hayes-Roth, B., An Architecture for Adaptive Intelligent Systems, Artificial Intelligence, 72 (12): 329-365, 1995.

[3] Anderson, J. R. and M. Matessa, In Oaksford, M. and Chater, N. (Eds.), The Rational Analysis of Categorization and the ACT-R Architecture, Oxford University Press, 1998.

[4] Wahl, S. and H. Spada, Children's Reasoning about Intentions, Beliefs, and Behavior. Cognitive Science Quarterly, 1: 5-34, 2000.

[5] Sloman, A., Varieties of Affect and the CogAff Architecture Schema. School of Computer Science, University of Birmingham, UK, 2001.

[6] Singh, P. and M. Minsky, An Architecture of Diversity for Common sense Reasoning. IBM Systems Journal, 41 (3): 530-539, 2002.

[7] Newell, A., Unified Theories of Cognition, Harvard University Press, 1990.

[8] Norman, D.A., Twelve Issues for Cognitive Science, Cognitive Science, 4: 1-33, 1980.

[9] Davis, D. N. and Lewis, S. J., Computational Models of Emotion for Autonomy and Reasoning. Informatica (Special Edition on Perception and Emotion Based Reasoning), 27 (2):159-165, 2003.

[10] Lewis, S.C., Computational Models of Emotion and Affect. Ph.D. Thesis, Department of Computer Science, University of Hull, UK, 2004.

[11] Davis, D. N., Cognitive Architectures for Affect and Motivation, Cognitive Computation, 2: 199-216, 2010.

[12] Piaget, J., Judgment and Reasoning in the Child, London: Routledge & Kegan Paul, 1928.

[13] Sloman, A., How many separately evolved emotional beasties live within us. Emotions in Humans and Artifacts, 35-115, MIT Press, 2003.

[14] Truran, K., Children's Use of a Representative Heuristic, In Proceedings of the 20th International Conference for the Psychology of Mathematics Education, 4: 339-346, 1996.

[15] Stachniss, C. and Burgard, W., Mobile robot mapping and localization in non-static environments. In Proceedings of the Twentieth National Conference on Artificial Intelligence, pages 1324-1329, 2005.

[16] Wolf, D. and Sukhatme, S. Mobile robot simultaneous localization and mapping in dynamic environments. Autonomous Robots, 19 (1): 53-65, 2005.

[17] Fiala, M. and Basu, A., Robot navigation using panoramic tracking. Pattern Recognition, 37 (11): 2195-2215, 2004.

[18] Rodriguez, M. Iglesias, R. Regueiro, C.V. Correa, J. and Barro, S., Autonomous and fast robot learning through motivation, Robotics and Autonomous Systems, 55 (9), p.735-740, Sep 2007.

[19] Stoytchev, A. Some Basic Principles of Developmental Robotics, IEEE Transactions on Autonomous Mental Development, 1 (2): 122-130, 2009.

[20] Davis, D. N. and Gwatkin, J., robo-CAMAL: A BDI Motivational Robot, Journal of Behavioral Robotics, 1 (2), Pages 116-129, July 2010.

[21] Arkin, R., Integrating behavioural, perceptual and world knowledge in reactive navigation. Robotics and Autonomous Systems, 6: 105-122, 1990.

[22] Stoytchev, A. and Arkin, R., Incorporating Motivation in a Hybrid Robot Architecture, Journal of Advanced Computational Intelligence and Intelligent Informatics, 8 (3): 269-274, May 2004.

[23] Georgeff, M. Pell, B. Pollack, M. Tambe, M. and Wooldridge, M., The belief-desire-intention model of agency. In J Muller, M Singh, and A Rao, editors, Intelligent Agents V: Agent Theories, Architectures and Languages, pp. 1-10. Springer-Verlag, 1999.

[24] Bartsch, K, and Wellman, H, Young children's attribution of action to beliefs and desires. Child Development, 60:946-964, 1989.

[25] Sloman, A., The Mind as a Control System. In C. Hookway and D. Peterson (Eds.) Philosophy and the Cognitive Sciences, 69-110, Cambridge University Press, 1993.

[26] Duffy, E, An Explanation of Emotional Phenomena without the Use of the Concept Emotion. The Journal of General Psychology, 25: 283-293. 1941.

[27] Simon, H. A, Motivational and Emotional Controls of Cognition, Psychological Review. 74 (1): 29-39, 1967.

[28] Sloman, A, Motives, Mechanisms, and Emotions. Cognition and Emotion, 1 (3): 217-234, 1987.

[29] Beck, R, Motivation: Theories and Principle, Pearson Education. 2000.

[30] Corkill, D.,D., Blackboard Systems. *AI Expert,* 6 (9): 40-47, 1991.

[31] Baars, B. J., In the Theatre of the Consciousness: The Workspace of the Mind, Oxford University Press, 1997.

[32] Davis, D..N. and Vijayakumar M V, A "Society of Mind," Cognitive Architecture based on the Principles of Artificial Economics International Journal of Artificial Life Research, Volume 1(1): 51-71, January-March 2010.

[33] Active Media. ARIA Reference Manual. ActiveMedia Robotics. 2010.

[34] Barsalou, L. W., Perceptual Symbol Systems, Behavioral and Brain Sciences, 22: 577-660, 1999.

[35] GC5 (2011) UK Computing Research Committee's Grand Challenge No. 5: Architecture of Brain and Mind, University of Birmingham website, http://www.cs.bham.ac.uk/research/projects/cosy/papers/aaai06-gc5.pdf.

[36] Langley, P., Laird, J. E. and Rogers, S., Cognitive Architectures: Research Issues and Challenges. Cognitive Systems Research, 10: 141-160, 2009.

[37] Chater, N., Tenenbaum, J.B. and Yuille, A. Probabilistic Models of Cognition: Conceptual Foundations. Trends in Cognitive Sciences, 10 (7): 287-291, 2006.

[38] Oaksford, M. and Chater, N. Précis of Bayesian Rationality: The Probabilistic Approach to Human Reasoning. Behavioural and Brain Sciences, 32: 69-120, 2009.

[39] Gibson, J. J., The Ecological Approach to Visual Perception. Houghton Mifflin, 1979.

DOI: 10.5963/IJCSAI0203001