Intelligent Motor Decision: From Selective Attention to a Bayesian World Model

Zenon Mathews, Sergi Bermúdez i Badia, and Paul F.M.J. Verschure, *Member, IEEE*

Abstract— **Intelligent sensor/motor allocation is gaining in importance in many areas of robotics and autonomous systems. It allows the autonomous entity to allocate its resources for solving the currently most critical task depending on the entity's current state, its sensory input and its acquired knowledge of the world. Such architectures which support dynamic motor allocation are invaluable for systems with limited resources. Biological systems also build and maintain a world-model to enable intelligent motor decision making. Based on recent advances in attention research and psychophysiology we propose a general purpose push-pull selective attention mechanism for building a world model and intelligent motor action control. We implement and test an architecture called A-BID, which is guided by a neural network implementation of a selective attention mechanism that is used to build a probabilistic world model. Using A-BID, the system performs at each time step the action that is optimal in the Bayesian sense.**

Index Terms— **Attention, autonomous systems, sensor fusion**

I. Introduction

Selective attention has recently become a key conceptual inspiration from neurobiology that motivates design of powerful information processing systems [3,4,7]. In such frameworks, an attentional window or "spotlight of attention" usually contains a subset of the sensory data perceived by the information processing system [6]. This "spotlight of attention" is forwarded to higher-order processes that plan and trigger responses of the system. Such architectures are vital for systems with limited resources. Often such systems are also designed to learn by acting, under utilization of attentional strategies [7]. Nevertheless, it is often unclear how the data-flow between the different components of such complex systems should ideally be modelled.

Based on psychophysiological studies a push-pull mechanism for selective attention in the human extras-triate cortex has recently been proposed [1]. Such a mechanism allows for optimal data-flow between the different subsystems and allows load balancing. Here we propose a general purpose push-pull mechanism of selective attention for a system that learns through action. Attentional spotlight is generated by a winner take all neural network and this is used as a goal-oriented attentional behavior. In our experiments we explore how, for a given complex task, attention guided world model is used to perform actions that are optimal in the Bayesian sense. We propose a general-purpose architecture named A-BID (Attention-Bayesian-Intelligent-Decision) which employs biologically inspired mechanisms enabling it to alternate between volitional, top-down and reflex level actions to maintain coherence of action. Our model was inspired by the distributed adaptive control (DAC) architecture proposed earlier for controlling behavioral systems that learn to act optimally to solve a given task [11,12]. Here we propose how DAC could be improved using attentional modulation of sensory data. In the next sections we lay the mathematical foundations for A-BID's data association, world model building and decision making and introduce our neural network implementation for selective attention. Finally we discuss a concrete robot simulation task to validate the systems performance.

II. Managing Bottom-Up Multimodal Sensory Data

The question we address in this section is how an autonomous entity manages the amount of multimodal sensory information it receives continuously. In the following, *stimulus* refers to a single modal observation (or data unit) and *target* means a well-defined physical object that also exists in the same space as the autonomous entity. Targets are perceived by the autonomous entity through the multimodal stimuli they evoke. The stimuli first have to be associated to existing targets, or if the stimuli is spatiotemporally distant from existing targets new targets have to be created. In this section we discuss the so called data association (or data alignment) problem.

A. Joint Probabilistic Data Association

Joint probabilistic data association (JPDA) is a powerful tool for solving data association problems, which arises in many applications such as computer vision, surveillance, mobile robots etc. JPDA is a single-scan approximation to the optimal Bayesian filter, which as-

Zenon Mathews is with the SPECS laboratory at Institut Universitari de l'Audiovisual (IUA) of the technology department of Universitat Pompeu Fabra (UPF) in Barcelona, Spain (telephone: +34 93 542 22 01, email: zenon.mathews@upf.edu)

Sergi Bermúdez i Badia is with the SPECS laboratory at Institut Universitari de l'Audiovisual (IUA) of the technology department of Universitat Pompeu Fabra (UPF) in Barcelona, Spain (telephone: +34 93 542 13 69, email: sergi.bermudez@upf.edu)

Paul F.M.J. Verschure is with the SPECS laboratory at Institut Universitari de l'Audiovisual (IUA) of the technology department of Universitat Pompeu Fabra (UPF) in Barcelona, Spain and with Institució Catalana de Recerca i Estudis Avançats (ICREA) Barcelona, Spain. (telephone: +34 93 542 2163, email: paul.verschure@iua.upf.edu)

sociates latest observations to known targets sequentially. JPDA enumerates all possible associations between observations and targets at each time step and computes the association probabilities β_{jk}, which is the probability that the j-th observation originated from the k-th target. Given such association probabilities, the target state is estimated by Kalman filtering [9] and this conditional expectation of the state is weighed by the association probability. Let x_t^k indicate the state of target k at time step t, ω_{jk} the association event where the observation j is associated to target k and $Y_{1:t}$ stays for all the observations from time step 1 to time step t. Then the state of the target can be estimated as

$$E(x_t^k|Y_{1:t}) = \sum_{\omega} E(x_t^k|\omega, Y_{1:t}) P(\omega|Y_{1:t}) \qquad (1)$$

$$= \sum_{j} E(x_t^k|\omega_{jk}, Y_{1:t}) P(\omega_{jk}|Y_{1:t}) \qquad (2)$$

where ω_{jk} denotes the association event where observation j is associated to target k and ω_{0k} denotes the event that no observation is associated to target k. Therefore the event association probability is

$$\beta_{jk} = P(\omega_{jk}|Y_{1:t}) \qquad (3)$$

JPDA computes a validation gate for each target using the kalman innovation of new observations. It only considers observations inside the validation gate for each target. For further mathematical details of JPDA see [9].

The computation of β_{jk} requires a summation over the posteriors and its exact calculation is NP-hard and is the major drawback of JPDA [5]. The number of association events rise exponentially in relation to the number of observations. We therefore implemented a Markov Chain Monte Carlo method to compute β_{jk} in polynomial time similar to the proposal by Oh and Sastry in [15].

B. Markov Chain Monte Carlo Computation of β_{jk}

In our system, the Markov Chain Monte Carlo (MCMC) method is used to estimate the association event probabilities β_{jk} in polynomial time and with good stability as shown in [15]. For this the problem is reformulated as a bipartite graph. Consider the bipartite graph $G = (U, V, E)$, where U is a vertex set of predicted observations, V is vertex set of observations and E is mapping predicted observations (or in other words a target) to an observation. We thereby only consider feasible mappings, i.e. the ones that respect the validation gate criteria for the JPDA. The algorithm starts with one such feasible mapping and a Markov chain is generated. For details of the MCMC approximation of β_{jk}, its convergence and stability see Oh and Sastry [15]. The polynomial time complexity with respect to the number of targets allows MCMC to be used for computing β_{jk} in real-time.

III. Goal Oriented Selective Attention

Biological nervous systems are still intriguingly superior to what can technically be implemented today for sensory data processing. Recent research has shown that the retina transmits between one and ten million bits per second, which is about the same rate as an Ethernet connection could support, to the brain [13]. Here we explore how attentional selection can add functional advantages to behavioral systems that deal with large amounts of sensory data. In particular, we consider here goal-oriented top-down selective attention as an information bottleneck that filters the most relevant sensory data, depending on the current task of the system [14], [1]. Such an information bottleneck, that changes dynamically with the system task, is critical for the survival of biological organisms as the incoming sensory data clearly overwhelms the available limited computational resources. In the next subsections we discuss our neural network implementation of such a goal-oriented selective attention mechanism and a load-dependend push-pull mechanism for optimal load balancing.

A. Attentional Spotlight Generation

For the implementation of the selective attention mechanism we use the IQR system for distributed large-scale real-time real-world neuronal simulations [18]. IQR allows implementing large neural networks for real-time applications and interfacing them to real world devices. As suggested by Itti and Koch [6], we implemented a set of neuronal feature filters and excitatory, inhibitory and time-delayed connections between them for saliency point computations. The selection of the feature filters depends on the current state of the system. E.g. if the system is running out of power, the feature filters for the charger have stronger excitatory influence on the salience computation. In the experiment section examples of such feature filters are discussed in more detail. This computation delivers goal-dependend salient target locations which are a subset of the total number of targets the data association mechanism has computed before (see figure 1).

B. Load-Dependend Push-Pull

Recent psychophysiological research suggests that the selective attention is load-dependend, i.e. how many unattended stimuli are processed depends on the degree to which attentional resources are engaged by an attended stimulus [1]. This provides evidence for a load-dependent push-pull protocol of selective attention operating at intermediate processing stages of the sensory-data. Such a push-pull protocol has behavioral effects for an autonomous system: when the attentional load is low the system can allocate motor and computational resources for unattended targets. Our architecture makes use of this load-dependend push-pull mechanism and this allows the acting system to switch between volitional, reflexive and explorative behaviors.

IV. World Model and Bayesian Decision Making

The world model of an attention-guided behaving system should ideally consist of the targets it attends to, but also if possible the unattended targets. Such a world model, or dynamic memory, allows the system to plan its actions depending on the top-down attentional load
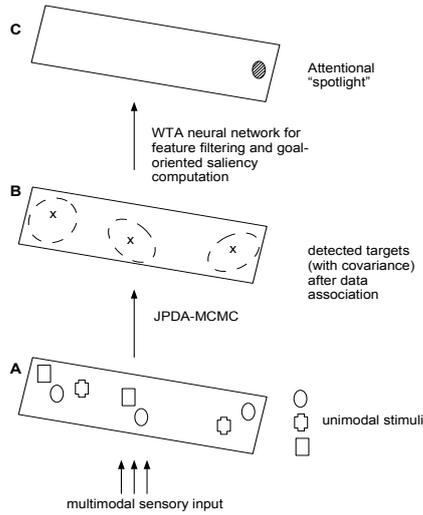
Fig. 1. Goal-oriented saliency computation from multi-modal sensory input: The multimodal sensory input (**A**) to the JPDA-MCMC algorithm is associated to existing targets (**B**). If the stimuli are distant enough from already existing targets, new targets are created. The goal-dependend saliency computation filters output the most salient targets depending on the current state of the system and the task at hand (**C**) .

and the bottom-up sensory input. In this section we discuss the building, maintenance and the use of such a world model for decision making.

A. A Dynamic Transient World Model

The data association and the attentional mechanisms deliver constant input to the world model. Our world model contains the spatial and temporal information of a total set of targets with the attended ones being represented more relevantly than the unattended ones. We write Θ_s^t for relevance of a certain target s at time t, and are interested in the following conditional probability:

$$P(\Theta_s^t|F_s^t(\Theta_s^{t-1})A_t(s)) \quad (4)$$

where $F_s^t(\Theta_s^{t-1})$ and $A_t(s)$ are two time-dependend functions which weigh the target s. For example $F_s^t(\Theta_s^{t-1})$ evaluates the spatial proximity of the target if there is at least one onset stimulus associated to this target and decays the current weight of the target otherwise. Whereas $A_t(s)$ evaluates the goal-dependend attentional saliency of this target. By computing the joint distribution of these relevance probabilities for all targets s the system can perform the motor action appropriate for the most relevant targets. The following subsection elaborates the update of these relevance probabilities. In the experiments section a concrete example for the world model and the Bayesian decision making is discussed.

B. Optimal Bayesian Decision Making

Let us assume that we can compute relevance probabilities of individual targets as shown above in eq. 4. Given these individual target relevances we are interested in the fused relevance distribution:

$$P(\Theta^t|F^t(\Theta^{t-1})A_t) \quad (5)$$

We express this probability as the normalized sum of probabilities of individual relevances:

$$P(\Theta^t|F^t(\Theta^{t-1})A_t) = \sum_s P(s)P(\Theta_s^t|F_s^t(\Theta_s^{t-1})A_t(s)S) \quad (6)$$

where random variable $S \in 1...n$, n being the number of targets and $P(s)$ indicates the probability of this target. As $P(s)$ is uniformly distributed over all targets this allows for normalization.

V. A-BID Architecture

The above sections discussed the individual computational modules of the A-BID architecture. This section elaborates on a general overview with all its subparts. The bottom-up multimodal sensory data is continuously *pushed* by the individual sensors to the data association mechanism, which aligns the multimodal stimuli to already existing targets or creates new targets. The result of this is forwarded to the world-model but also to the saliency computation module. In parallel the goal-oriented attentional spotlight generation modulates the relevance of target representations in the world model so that depending on the current task the representation of relevant targets are enhanced. In the world model the relevance of the individual targets are represented probabilistically. The Bayesian decision making module operates on this world model and generates motor actions which are optimal in the Bayesian sense.
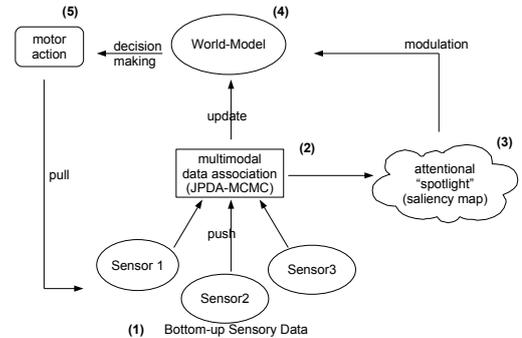

Fig. 2. A-BID Architecture and the Push-Pull Data Flow: 1) The bottom-up multimodal sensory data are *pushed* as they become available. 2) The data association mechanism aligns the multimodal stimuli to already existing targets or creates new targets. 3) The goal-oriented attentional spotlight generation mechanism modulates the relevance of target representations in the world model so that depending on the current task the representation of relevant targets are enhanced. 4) The world model is a probabilistic representation of the relevance of the targets at hand. The Bayesian decision making grounds on this world model and generates motor actions that are optimal in the Bayesian sense.

VI. Methods

To test the above discussed architecture we consider the following robot swarm scenario. A swarm of robots

are on a common mission in a given environment. One of the robots named A-BID is involved in the specific task of aiding expired (i.e. broken-down or out of charge) agents. This means that A-BID first has to localize the expired agents using its sensors and approach them for repair or recharge. A-BID is equipped with a limited number of distance-measurement sensors like sonar and laser range scanners. With the range-scanners A-BID has to scan the environment and recognize the agents to accomplish the given task. From time to time, A-BID has to go back to the base station to recharge itself. Solving this multiple goal task involves multimodal data association, goal-driven selective attention generation and maintaining a world model, which is used to compute the optimal action in the Bayesian sense. We simulated this real-world scenario and used the A-BID system for solving the given task (figure 3).
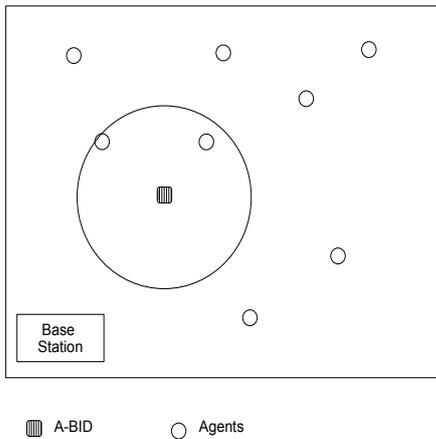

Fig. 3. A-BID experiment setup: A-BID can only perceive the world inside the indicated circle as the range of its sensors is limited. Nevertheless, A-BID's world model contains agents outside of its sensory range, if it had perceived them in the past. The probability about the certainty of the state of such targets outside A-BID's sensory range will constantly decay with time, which allows A-BID to slowly forget what it had seen in the past. All the involved entities are constantly in motion. A-BID has the task to locate non-moving agents and approach them to recharge and/or repair.

A. A-BID Experiment: World Model and Bayesian Decision

We simulate the data from different range sensors and use this as multimodal stimuli to the JPDA-MCMC algorithm. This creates targets of which the relevance probability has to be computed. In the following we derive the necessary concrete equations for the general ones shown in 4.

The general forms of equations 4, 5 and 6 have to be instantiated for this concrete experiment as in the following. Let us assume that the motor action consists simply of choosing a direction of motion $\gamma \in 0..360$ and a distance $\psi \in 1..10$. The best action is then chosen in the direction γ of the most relevant target at distance ψ in the world-model. We are interested in computing the most relevant direction of motion γ and distance ψ. Therefore we are interested in the probability:

$$P(\gamma\psi|F_s^t(\Theta_s^{t-1})A^t(s))$$

As F is a function of distance d and time t we can express F as the known quantities for a single target. These quantities are, the distance d to A-BID, the time t since the time of a previous stimulus associated to this target, the orientation γ_i relative to A-BID and the attentional weights a_i for each target. This lets us formalize the above probability as:

$$P(\gamma\psi|d_1,...d_nt_1,...t_n\gamma_1,...\gamma_na_1,..,a_n) \quad (7)$$

First we consider the conditional probability 7 as if there were only one target i and without attentional inputs a_i. Assuming conditional independence for angle and distance domains we do the following decomposition:

$$P(\gamma\psi|d_it_i\gamma_i) = P(\gamma|d_it_i\gamma_i)P(\psi|d_it_i\gamma_i)P(t_id_i\gamma_i)$$

The probability distributions $P(\gamma|d_it_i\gamma_i)$ and $P(\psi|d_it_i\gamma_i)$ are Gaussian distributions with the following means and standard deviations:

$$P(\gamma|d_it_i\gamma_i) = \mathcal{N}(\gamma_i, \frac{d_it_i}{c_1}) \quad (8)$$

$$P(\psi|d_it_i\gamma_i) = \mathcal{N}(c_2d_i, c_3t_id_i) \quad (9)$$

and we assume the uniform distribution

$$P(t_id_i\gamma_i) = \mathcal{U} \quad (10)$$

where c_1, c_2 and c_3 are constants.

The semantics of the above equations are the following:

Equation 8 : this Gaussian is centered on the angle γ_i at which the agent i is located. The standard deviation is a function of time t_i at which this agent was last perceived and the distance d_i at which this agent is. This allows A-BID to gradually forget (time decay) what it had seen in the past, as past information is always prone to change in a dynamic world.

Equation 9 : analogously for the distance domain, the Gaussian is centered on the distance d_i of the agent and the standard deviation is again a function of time t_i, allowing a time decay.

Equation 10 : this joint probability is uniform as we do not have any prior information about possible correlations between those random variables.

Now to consider the relevances of all the targets for the computation of the total relevance as shown in equation 6, which we bring back the attentional components a_i and consider the following conditional probability distribution:

$$P(\gamma\psi|d_1,...d_nt_1,...t_n\gamma_1,...\gamma_na_1,..,a_n) = \quad (11)$$

$$\sum_i \frac{a_i}{a_{tot}}P(\gamma|d_it_i\gamma_i)P(\psi|d_it_i\gamma_i)P(t_id_i\gamma_i) \quad (12)$$

where a_{tot} is the sum of all attentional components a_i.

This means that the attentional components a_i weigh the shares of the individual agents to the joint conditional probability distribution. In other words, attention modulates the world model, which is expressed as a probability distribution that changes in each step with the sensory input.

B. A-BID Experiment: Selective Attention Computation

The robot A-BID can have several goals at any given time. When it is not running out of charge its goal is to localize other robots which are not moving and approach them for aid. On the other hand, when A-BID is running out of power, its goal is to get back to the base-station for recharging itself. Therefore the top-down goal-oriented saliency computation takes the current state of A-BID into account. Different feature detectors were implemented to code for selected features such as motion, density of targets etc. The saliency computation basically delivers the a_i attentional component for the probability computation in equation 12. Figure 4 describes the neural network implementation for saliency computation, done using the large-scale neuronal simulator [18].
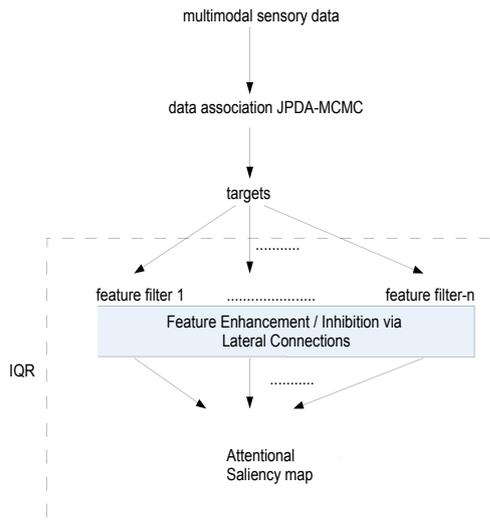

Fig. 4. A-BID Neural Network Implementation for Attentional Spotlight Computation: The large-scale neuronal simulator IQR is used to compute the saliency map using the feature detectors. The target data (after data association using JPDA-MCMC) is filtered using various feature filters. Some feature filters like, motion, density, loneliness etc. are used here to compute salient points.

For the Bayesian probability computations, the ProBT library was employed [19]. All implementations were done in a C++ Linux environment on a Pentium M 1.7 GHz processor.

VII. Results

We used 10 agents who move around the experimental field of a 600 by 600 pixels window. They start with a maximum speed and maximum energy but slow down as the energy drops. The energy-drop is proportional to the covered distance. Their direction of motion is arbitrary but always inside the field. A-BID always starts from the base station and alternates between explorational and exploitational time slots. During exploration it moves about the field randomly to detect the agents. Thereby the multimodal sensor fusion and attentional saliency computation delivers input to update A-BID's world model. During the exploitational time slot, A-BID performs intelligent motor action as described in the methods section. We compare the performance of A-BID in two categories: use the world model in the exploitation phase and in the lack of it. When the world model is not used to compute an intelligent action, A-BID is in constant exploration. For each category 5 trials each with 5000 time-steps were carried out. Figure 5 illustrates the performance during the simulation. The probabilistic world model computed for this instance is shown in figure 6.
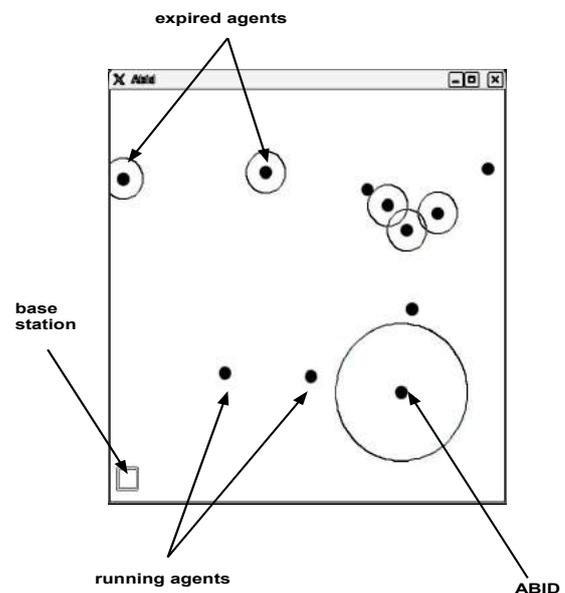

Fig. 5. Snapshot from A-BID Experiment: The bigger circle surrounding A-BID shows the sensory range of A-BID. Everything outside this range is not perceived by A-BID. The smaller circles surrounding some of the agents signify that the corresponding agents are expired. See figure 6 for the computed probabilistic world model for this situation.

Figure 5 shows a snapshot of the A-BID simulation GUI. The probabilistic world model for this particular situation is shown in figure 6. This Bayesian world model allows A-BID robot to choose the probabilistically best motor action.

We look at the number of recharged agents during each trial and also at the total expiry time of all agents together in each run and observe a significant improvement when A-BID uses the probabilistic world model and motor action selection (WM indicates the use of world model and non-WM indicates the use of a reactive system that explores the robot arena without a world-model or attentional mechanisms)). (See figures 7 and 8).

VIII. Conclusion and Future Work

Here we introduced our first mathematical framework and computational architecture of a Bayesian world
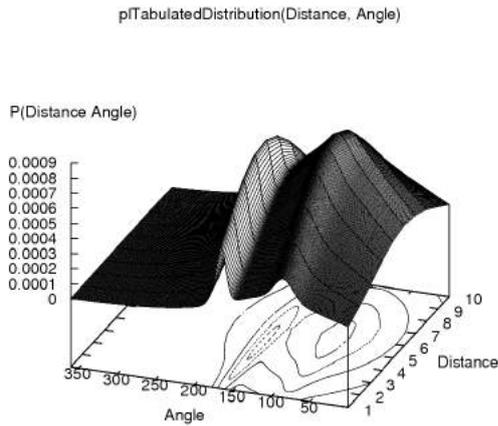
Fig. 6.   Example of A-BID world model probability manifold: Figure shows the probability distribution computed for the instance shown in figure 5. Angles range from 0..360 in anti clockwise direction and distances range from 1 to 10.. As A-BID had "seen" before that there are some slow agents (higher saliency) along the 90 degree direction, the probabilities for those areas are higher. The peak is at 85 degrees and distance 6, which corresponds to the positions of the clutter of expired agents in figure 5. This world model of A-BID suggests the most probable action as the one that leads to the agents, which were perceived to be running slow in the past. This probability distribution is computed at each time step before A-BID takes an intelligent motor action decision.
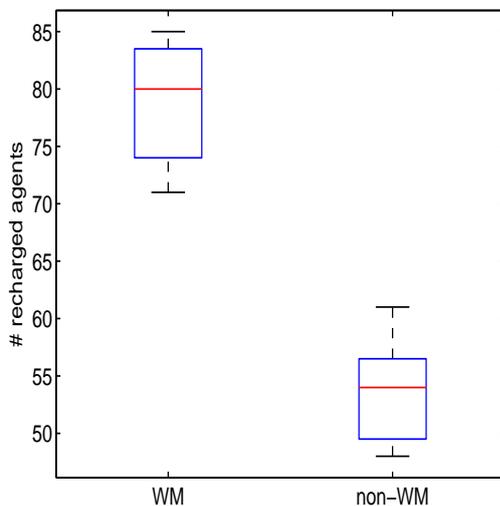

Fig. 7.   Number of recharged agents were significantly higher (left boxplot) when using the A-BID world-model architecture than when not using it (right boxplot).

model, which uses multimodal data association and selective attention mechanisms to enable intelligent motor action selection. Our probabilistic model was validated using a multi-goal robotic swarm simulation. In future we plan to validate our model on real robots for solving real world tasks under limited resources and time constraints.
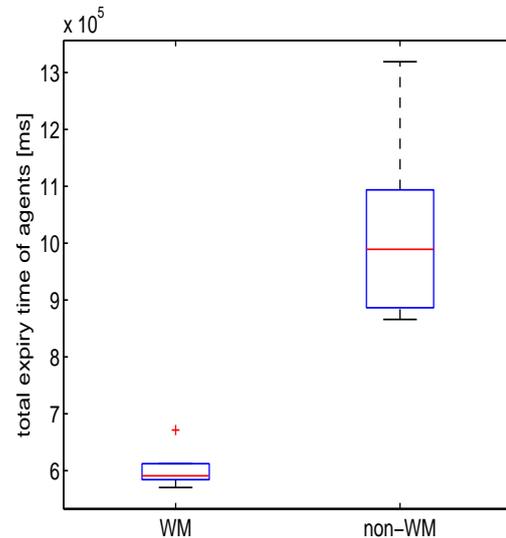

Fig. 8.   The total expiry time of the agents, i.e. their inactive time as they had run out of charge, is significantly less when using the A-BID world-model architecture (left boxplot) than when not using it (right boxplot).

REFERENCES

[1] M. A. Pinsk, G.M. Doniger and S. Kastner, *Push-pull Mechanism of Selective Attention in Human Extrastriate Cortex*, Journal of Neurophysiology 92, 2004.
[2] S. Kastner and L.G. Ungerleider, *Mechanisms of Visual Attention in the Human Cortex*, Annu. Rev. Neuroscience 2000.
[3] W.J. Won, S.W. Ban and M. Lee, *Real Time Implementation of a Selective Attention Model for the Intelligent Robot with Autonomous Mental Development*, IEEE ISIE 2005.
[4] Y. Jiang, N. Xiao, L. Zhang, *Towards an Efficient Contextual Perception for Humanoid Robot: A Selective Attention-based Approach*, 6th World Congress on Intelligent Control and Automation 2006.
[5] J. Collins, J. Uhlmann, *Efficient gating in data association with multivariate distributed states*, IEEE Trans. Aerospace and Electronic Systems, vol. 28, no. 3, pp. 909-916, 1992
[6] L. Itti, C. Koch, *Computational Modeling of Visual Attention*, Nature Reviews in Neuroscience 2 (2001)
[7] G. Billock, C. Koch, D. Psaltis, *Selective Attention as an Optimal Computational Strategy*, Neurobiology of Attention. p. 18..23 Elsevier Academic Press 2005
[8] T.J. Anastasio, P.E. Patton, K. Belkacem-Boussaid, *Using Bayes' Rules to Model Multisensory Enhancement in the Superior Colliculus* Neural Computation 12, 1165-1187, 2000
[9] Y. Bar-Shalom, *Tracking and Data Association*, Academic Press Professional Inc. San Diego CA USA 1987
[10] K. Eng, R.J. Douglas, P.F.M.J. Verschure, *An interactive space that learns to influence human behavior*, IEEE Transactions on Systems, Man, and Cybernetics, 2005, 35:(1) 66-77.
[11] P.F.M.J. Verschure, T. Voegtlin, R.J. Douglas, *Environmentally mediated synergy between perception and behaviour in mobile robots*, Nature 2003, 425: 620-624.
[12] P.F.M.J. Verschure, P. Althaus, *A real-world rational agent: Unifying old and new AI*, Cognitive Science 2003, 27: 561-590.
[13] C. Koch, J. McLean, R. Segev, M.A. Freed, M.J. Berry II, V. Balasubramanian, P. Sterling, *How much the eye tells the brain*, 2006 Curr Biol 16:1428-1434
[14] N. Navalpakkam, L. Itti, *Search Goal Tunes Visual Features Optimally*, Neuron 53, 605-617, 2007 Elsevier Inc.
[15] S. Oh, S. Sastry, *A Polynomial-Time Approximation Algorithm for Joint Probabilistic Data Association*, in Proc. of the American Control Conference (ACC), Portland, OR, June 2005
[16] Open CV, *Intel Open Source Computer Vision Library*
[17] B.E. Stein, M.A. Meredith, *The Merging of the Senses*, MIT Press Cambridge MA 1993
[18] U. Bernardet, M.J. Blanchard, P.F.M.J. Verschure, *IQR: A distributed system for real-time real-world neuronal simulation*, Neurocomputing 2002, 44-46: 1043-1048
[19] *www.bayesian-programming.org*
[20] Z. Mathews, S. Bermdez i Badia, P.F.M.J. Verschure, *A Novel Brain-Based Approach for Multi-Modal Multi-Target Tracking in a Mixed Reality Space*, Proceedings of 4th IN-TUITION International Conference and Workshop on Virtual Reality 2007, 4-5 October Athens, Greece.