The Fusion of Multiple Sources of Information in the Organization of Goal-Oriented Behavior: Spatial Attention versus Integration

Milanka Ringwald* Paul F.M.J. Verschure†
* *Institute of Neuroinformatics, UZH/ETH Zurich, Zurich, Switzerland*
† *ICREA and Department of Technology, University Pompeu Fabra, Barcelona, Spain*

Abstract— **In the fields of neuroscience, psychology and robotics, an important question is how to establish a unified system that will autonomously acquire both its input state space and optimal goal-oriented action policies in unknown environments. An important requirement for a such a system is to understand how multiple sources of sensory information can be integrated to support autonomous behavior. So far, Distributed Adaptive Control (DAC), a self-contained neuronal system, used only egocentric cues to achieve goal-directed behavior in a foraging task. However, implicitly acquired navigation strategies are not well exploited. In this paper, we evaluate the hypothesis that learned ego-centrically defined behavioral strategies can be improved by the integration of allocentric spatial information. Using an extension of the DAC architecture in the context of random foraging, we show that this integration can be rather seen as an instance of Bayesian inference as opposed to selective attention. We provide an extensive analysis of the architecture and compare its performance in the broader context using a known robotics algorithm. Our results further support the belief that a Bayesian framework can provide for a unified view on the organization of goal-oriented behavior.**

I. Introduction

An important requirement for the development of fully autonomous robots is a self-contained system supporting perception, cognition and behavior. Important aspects of such a system are both unsupervised learning of the input state space from multiple information sources and optimal policies of integration in goal-oriented behavior in unknown environments. For modeling of reasoning under uncertainty in robotics, we often encounter methods for multi-sensor fusion based on approximations of Bayesian inference [1] such as Dynamic Bayesian Networks (DBN) for goal-oriented action sequence selection [2, 3]. These models have been shown to work well for navigation within small-scale environments but they usually make very strong assumptions, i.e., having a predefined state space or globally defined frames of reference. Hence, although shown to be robust to the inherent uncertainty of their sensory inputs and/or the partially available information about the environment, they use global reference frames to define an action (e.g. go north, east, west, south) and to compensate for errors in state estimation such as those resulting from the integration of odometry information. This raises the question how we can achieve optimality in a self-contained architecture that does not require a priori global information. In this paper, we will present a robot based neural model, called Distributed Adaptive Control (DAC) [4, 5, 6, 7] and its extentsion, that shows that the perceptual and behavioral requirements of decision making can be met in the context of a Bayesian framework.

Support for the idea that the brain performs Bayesian integration comes from work on human decision making where it is shown that optimal predictions in everyday cognitive tasks and on the actions of an intentional agent, based on observing its behavior, can be explained in a Bayesian context [8, 9]. These and related observations have given rise to the field of neuroeconomics where the neuronal mechanism underlying optimal decision making are seen in the context of Bayesian and related methods [10]. A meaningful Bayesian prediction requires the existence of adequate prior knowledge and, in this sense, represents an hypothesis on the situation to which a subject is exposed. This suggests that the brain, as a result of having acquired correct priors either by phylogenetic or ontogenetic means, can produce meaningful predictions based on the statistics of the world, even in cases where the amount of sampled stimuli is small. Although Bayesian theory is very general and widely applicable to many problems in decision making, it provides only a very abstract description of the observed phenomena. For instance, the most elaborate and behaviorally validated models of human speech perception are defined as abstract functional models expressed in fuzzy logic [11, 12]. Hence, the neuronal mechanisms that could underlie Bayesian optimal decision making are still unknown.

The optimality of the behavior is not restricted solely to the human brain, it also has been assessed by behavioral experiments on animals [13]. For instance, rats placed in a radial arm maze, where different arms contain a varying amount of food pellets, develop an optimal foraging strategy in terms of travel time, probability of food occurrence and amount of food that adapts to changes in environment [14]. It has been shown that the adopted strategies are strongly influenced by the expected gain and its magnitude [15] and that these strategies maintain an optimal balance between exploration and exploitation [16]. Complementing the results of the behavioral studies, in the study of the neuronal substrate underlying spatial navigation, it is well accepted that the hippocampus plays a crucial role in the generation of contextual information [17]. The hippocampus contains neurons that show a response that is modulated by the location of the animal in its arena: i.e place fields. These place fields can be understood as a combination of view-dependent Gaussian shaped receptive fields [17, 18] and they are thought to provide an organism with information on its instantaneous allocentric spatial location.

It is currently not clear how this spatial information can be integrated with other sources of information in order to generate structured goal-oriented behavior. We study this question in the context of a virtual mobile robot that performs a random foraging task in an open arena. In particular, we compare integration scenarios based on either attentional selection or on probabilistic integration. We show that the Bayesian integration scheme display superior performance in the foraging task. Subsequently, we show that the extension of the DAC architecture based on augmented recall of behavioral sequences by a "spatial attention" signal or integration suggests local mechanisms that could facilitate the transformation of egocentrically defined actions into allocentric behavior.

II. Bayesian Inference within the DAC Framework

DAC evolved from a robot based neuronal model of classical and operant conditioning to a system that generates goal-oriented adaptive behavior that is derived from three tightly coupled layers of control: reactive, adaptive and contextual ([5]). The architecture of the contextual layer has analogs of the central components of a Bayesian analysis of the foraging task: goals, actions, hypotheses, observations, experience, prior probabilities and score function. By phrasing the foraging tasks performed with the DAC architecture in Bayesian terms it has been shown that the DAC architecture executes exactly those actions that are optimal in a Bayesian sense [6]. The following sections introduce the basic concepts of DAC and our extension of its contextual layer that support the integration of multiple information sources.

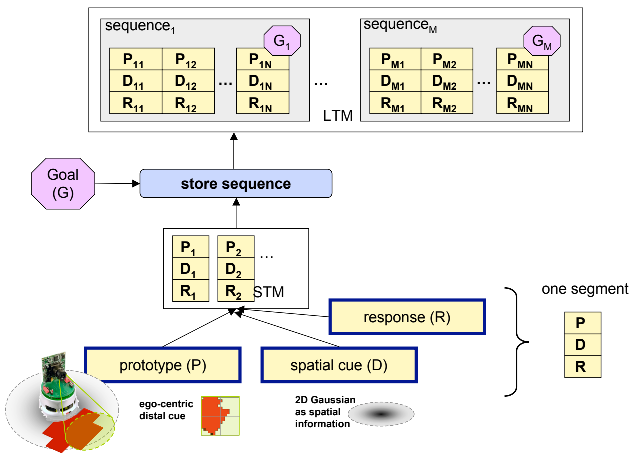

Fig. 1. DAC contextual layer: sequence learning process.

Reactive layer. The reactive layer comprises prewired reflexive relationships between simple sensory events and actions. In our experiments, these simple sensory events are defined as proximal cues, i.e., collision detection (punishment) and the intensity of a local target, i.e., light (reward) or a distance to the visual cue. If none of the cues are present, the reactive control layer generates exploratory behavior, i.e., translation.

Adaptive layer. At the level of adaptive control, distal input (colored patches) is used to form representations of cues that co-occur with the proximal cues (target/collision sensory input) that arise from the robot's actions. These representations are called prototypes. Over time, these learned prototypes progressively replace the purely reactive triggering of discrete actions. The local learning mechanism at this level dynamically adjusts a measure of discrepancy between expected and actual distal events. When the discrepancy falls below a transition threshold, the contextual control layer is enabled. This transition mechanism ensures that the prototypes used in the contextual layer are based on stable classifications constructed by the perceptual learning system of the adaptive layer.

Contextual layer. The contextual layer provides mechanisms for short- and long-term memory (STM/LTM). In the original DAC architecture only co-occurring prototypes and responses are stored in STM. In our extension, we additionally store allocentric spatial information, i.e., the position where an event has occurred. The general LTM concept still holds. A tuple containing a prototype, spatial information and a response represents a segment. Segments are stored in STM forming a sequence that conserves the order of their occurrences, see Figure 1. An STM sequence is stored in long term memory when a goal state, G, is reached, such as finding a target. The sensory content of LTM segments is continuously matched against interpreted sensory events, P, generated by the adaptive layer. The recall of behavioral actions from LTM is based on the matching of ongoing sensory events to those retained in memory. Action is estimated in a greedy manner as an average of the LTM actions weighted by the contribution of the matched LTM segments to which they are attached [5]. DAC estimates an action when new observation becomes available. A feedback mechanism favors segments following the selected one along the corresponding LTM sequence by temporarily increasing their weights.

We extended the architecture to provide support for our new concept of biasing an LTM sequence based on sequence fidelity. Sequence fidelity is illustrated on the Figure 3(a). It represents a mechanism for memory smoothing. In the original DAC architecture, the contribution of an LTM segment was provided by the value of its collector unit, c, where $\mathrm{c}_\mathrm{l}^\mathrm{k}$ of segment l of sequence k is defined as:

$$\mathrm{c}_\mathrm{l}^\mathrm{k} \quad = \quad \alpha(1-\mathrm{m}_\mathrm{l}^\mathrm{k}\mathrm{t}_\mathrm{l}^\mathrm{k}) \tag{1}$$

where $\mathrm{m}_\mathrm{l}^\mathrm{k}$ represents the distance of the prototype stored in the segment to the current one [5]. α represents the normalization constant that enforces the total sum of all LTM segment collector units to be one. The activation the trigger unit $\mathrm{t}_\mathrm{l}^\mathrm{k}$ of each segment falls back to its default value according to:

$$\mathrm{t}_\mathrm{l}^\mathrm{k}(\mathrm{t}+1) = \alpha_\mathrm{T} + (1-\alpha_\mathrm{T})\,\mathrm{t}_\mathrm{l}^\mathrm{k}(\mathrm{t}) \tag{2}$$

where $\alpha_\mathrm{T} \in [0;1]$. This trigger unit is used to model sequence fidelity as $(1-\mathrm{t}_\mathrm{l}^\mathrm{k})$. The collector unit of a segment is directly proportional to the conditional probability (likelihood) of observing the current prototype with respect to the information carried by the segment. The calculation of collector unit value precedes the action selection process. Its value represents the conditional probability $p(o|r)$ that the robot observes prototype o, given that a target will occur after n time steps executing the action r associated with the prototype of the LTM segment. This probability is greater than zero if there is a sequence in LTM, where the n-th to last segment has stored the action.

The justification of the probabilistic inference of the LTM segment weights, where the computation of the α value (Eq. 1) is simplified to the calculation of the normalization constant in a way it is proposed here, is given by the belief propagation method [19] and it is related to the action selection process as described in [2]. By incorporating another source of the sensory information in a Bayesian manner, the collector unit's value is given as:

$$c_l^k = \alpha(1 - m_l^k t_l^k)(1 - s_l^k t_l^k) \tag{3}$$

where m_l^k represents a gaussian with the distance measure between two prototypes (as defined in [5]) as its mean, and scale of spatial information as its standard deviation. Similarly, s_l^k represents a gaussian with mean equal to the Euclidian distance between the LTM segment and the current position of the robot, and the same scale of spatial information as its standard deviation.

III. Integration of Spatial Information

In this section, we present our two approaches for bimodal sensor fusion: *spatial attention* and *integration*. Both fusion mechanisms use spatial information to narrow down the number of matched segments from the LTM behavioral sequences to improve local awareness. The spatial attention approach improves local awareness by selecting only those "salient" cues that are coherent to spatial information. The spatial integration approach, in contrast, uses a Bayesian technique to bias cues.

In the original DAC architecture, an LTM segment contains an egocentric visual prototype and the action performed when the prototype was acquired. We augmented the LTM segments to hold also allocentric spatial information. We assume that allocentric spatial information is provided by the place cells of the hippocampus. As an abstract version of the place fields of hippocampal place cells, consistent with recent work [18], we used perfect 2D-Gaussian tuned place-fields.

Spatial Attention. In the spatial attention approach, the process of selecting LTM segments consists of two steps. In the first step, the similarity of the current visual cues to all LTM segments is measured. If the similarity reaches a given threshold, we say that the segment is matched. In the second step, the distance between the current 2D Gaussian place field and the place field of the matched segment is used to evaluate the quality of the segment. If the matched segment passes second step, its similarity with the visual cue will define the likelihood of its action. By this, possible ambiguity is reduced as only segments that are coherent with the spatial information are considered in the action selection process, see Figure 2(a).

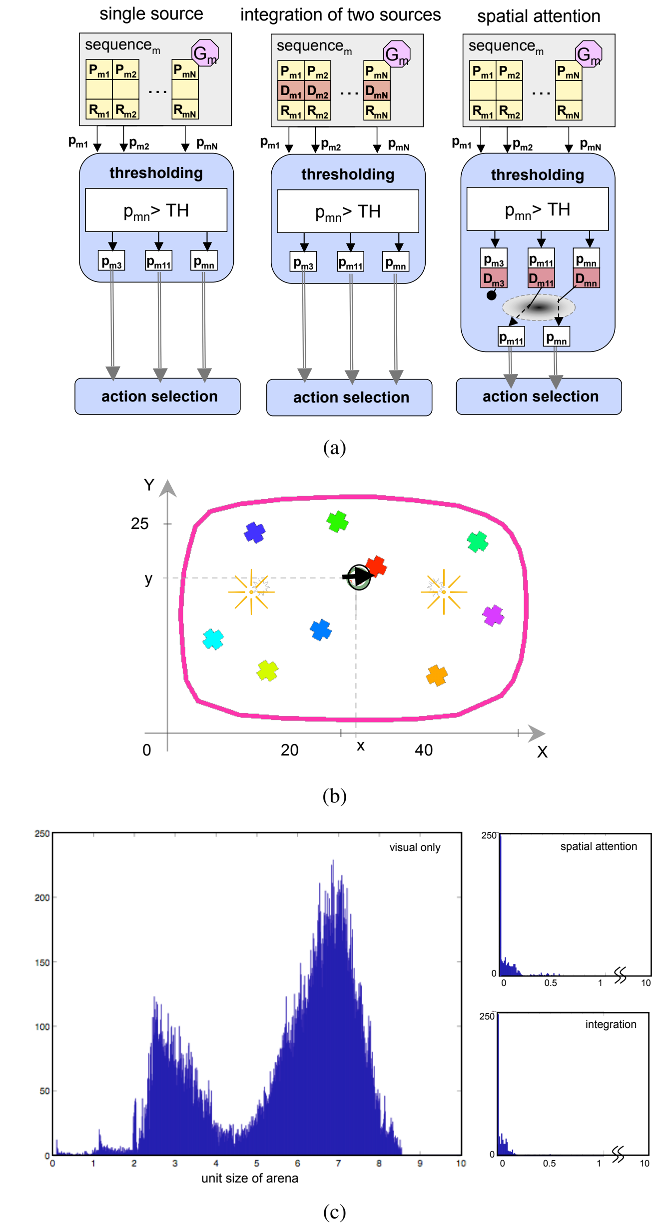

Fig. 2. a) Graphical representation of different scenarios of the sensory integration methods. b) Top-view on the arena. Two stars represent targets. Only one target is active at a time. As soon as robot finds active target, this target is deactivated. The size of the arena is 40 by 25 units, where the unit size equals to the size of a robot with respects to its camera properties. c) Each plot represents a histogram of the position estimation error during the recall experiment in the last 50000 cycles. Position estimation error of the vision only, spatial attention and integration method are shown. First peak in the position estimation error of the vision only approach is due to symmetry of the visual cue. Second peek is estimation error that occurs when the wall surrounding arena is in the sight of the robot.

If present, this ambiguity will lead to position estimate error, which will then have an effect on the performance of the robot, see Figure 2(c).

Spatial Integration. In the spatial integration approach, visual cues and spatial information are used to calculate the likelihood of actions stored in the memory. I.e., the value of the LTM segment collector unit with respect to the current observation and the robot position depends on both: the similarity between prototypes and the distance between the 2D Gaussian place fields, see Figure 2(a).

IV. Evaluation

We evaluate our extensions to the model together with additional control groups using a simulated robot in an open-arena foraging task, see Figure 2(b).

Firstly, we evaluate the performance of these two approaches with respect to two parameters, namely cue fidelity and sequence fidelity. We analyzed the possible cause for the improved performance. The first hypothesis is that the improvement is a result of better acquired sequences. The second hypothesis is that the acquired sequences are more efficiently exploited. To asses the quality of stored sequences, we employ the mutual information method [21], an information theoretic approach that enables us to quantify the amount of missing information in the LTM of one experiment with respect to the other experiment. We use logarithms to the base 2 such that the mutual information is given in the number of bits.

Secondly, we analyzed the trajectory of the robot during the experiment. We validate that the trajectory is caused by activity of the contextual layer and cannot be explained with the lower layers. For this, we recorded the LTM memory sequences for the three DAC levels separately: reactive layer, reactive and adaptive layers, and the complete hierarchy. Then, we modeled the LTM sequences of these levels as a hidden Markov model, HMM [20]. Having HMM models of the DAC levels enables us to generate optimal state sequences that we compare with the appropriate excerpts from the robot's trajectory which represent observation sequences. The HMM model consists of a set of states and transition probabilities between them. In addition, each state has an associated emission probability describing probability distribution for generating a symbol. The first step in creating a HMM from LTM is to define states and their emission probabilities, then capture transition probabilities between such states. LTM memory consists of a set of short sequences of ten segments capturing two to three different visual cues. The projection of LTM segments onto the arena surface reveals clusters centered around visual cues (symbols). We modeled each cluster as a separate state with Gaussian probability distribution. We generate these clusters in an unsupervised mannner by employing the Expectation-Maximization (EM) algorithm [22]. The EM is run until it converges. The transition probabilties between such clusters are then read out from the LTM sequences. We created observation sequences from the the robot's trajectory when memory was active. Each sequence consists of three visual cues. By having a set of observation sequences and their respective optimal state sequences, we can compare their distances. As a distance measure we used Kullback-Leibler divergence (relative entropy) as described in [23]. Using bit as a measurement unit for the two non-overlapping sequences we expect a KL divergence of two bits. Two correlated sequences have a KL divergence equal to zero. We assessed the likelihood that the observation sequence comes from each of these models and we looked at the correlation of these observation sequences with their "optimal" state sequences generated by each of DAC HMM models.

Finally, we evaluate the system with respect to the performance of its lower layers and in a broader context using a known robotics algorithm.

A. *Effects of Bimodal Sensory Integration*

To evaluate the performance of these two approaches, we performed four sets of experiments. The first uses the "spatial attention" while the second assesses the impact of the "spatial integration" method. We compared these results with two control experiments in which the recall of the behavioral sequences is based solely on either the spatial signal or the visual cues stored in LTM.

By comparing the target rates of the four types of experiments, we observed that both, the "visual cue" and the "spatial only" condition show a low performance, see Figure 3(c). We inspected the trajectories of the experiments runs of these two groups. For both, the robot falls into two behavioral patterns: moving along a stereotypical path including only one target or following the wall. Effect of such behavior is reflected on the content of the memory, i.g., the memory content of these two experimental groups shows significantly less information about the environment compared to the respective content in the best performing group (integration method with the spatial and memory smoothing scale: (0.5, 2.0))

The "integration" approach employs a Bayesian method to decide on the actions of the robot. This method shows the highest performance among all other, if the spatial scale is set low. Increase of the spatial and memory smoothing scales decreases the performance. However, if the increase of the uncertainty in the spatial information is counterbalanced by decrease in the scale of the memory smoothing, the performance is preserved.

In the "spatial attention" approach, only those cues and actions stored in LTM that have been stored in a region around the current position are selected. This method results in a low variance in memory content and performance across the given range of memory smoothing and spatial scale. As a result of the better exploration/exploitation strategy at the higher scales, it outperforms the "integration" method. While the weight of the action stored in the LTM segment is determined solely by the similarity of the disambiguated visual cue, these results further suggest that of the two sources of inputs, spatial and vision driven, the latter has a dominant role.

In summary, we showed that the integration of the two sensory inputs not only reduces ambiguities of the egocentric frame reference but, at the same time, it improves both: an exploratory behavior of the robot and its performance.

B. *Correlation of robots trajectory to LTM sequences*

The DAC system consists of three coupled layers. The lowest one provides prewired reflexive actions. Adaptive layer is associating distal cues with the low sensory inputs of the reactive layer, and over time it will overide the purely reactive response at the places where distal cues and low sensory inputs occur together. Contextual layer learns sequences of perception/action pairs that lead to goal states, either targets or collisions. We wanted to confirm that the performance

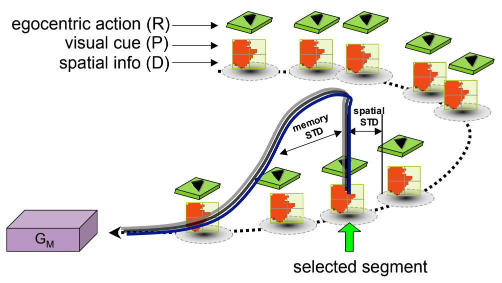

(a)

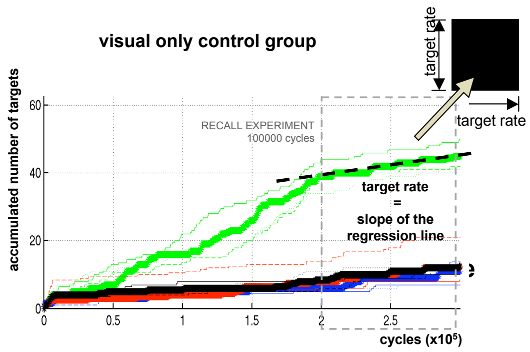

(b)

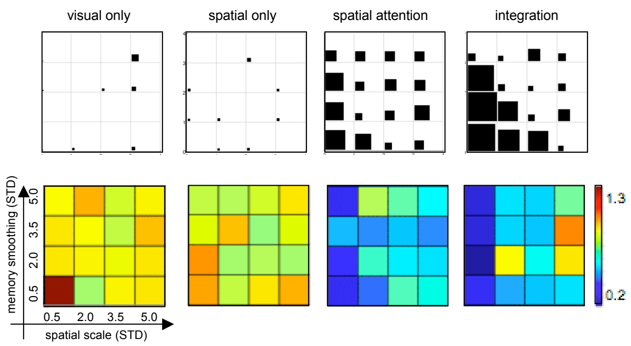

(c)

Fig. 3. Performance data. a) An unfolded LTM sequence. Each segment of the sequence contains an egocentric visual cue called prototype, a spatial information and an egocentric action. The scale of the spatial information represents the standard deviation of the 2D Gaussian, with the unit size equal to the size of the robot and it directly influences the cue fidelity. Another parameter is the scale of the memory smoothing. It is also represented with the standard deviation of the 2D Gaussian and is used to enhance sequence fidelity. b) We measured the mean number of accumulated targets over a fixed period of time (50 000 cycles). Each line represents an average over 40 experiments for a fixed set memory smoothing and spatial scale. The slope of the regression line is used as measure for the target rate. c) The plots show the target rate (upper plot) and information content(lower plot) for the four experimental groups by varying two parameters: cue and sequence fidelity. These two parameters are represented by spatial and memory STD respectively. The memory of the best performing group (integration method (0.5, 2.0)) is taken as the baseline to compare the information content of the memories stored by each experimental group. Color on the plots corresponds to the number of missing bits of information, with blue one being the minimum.

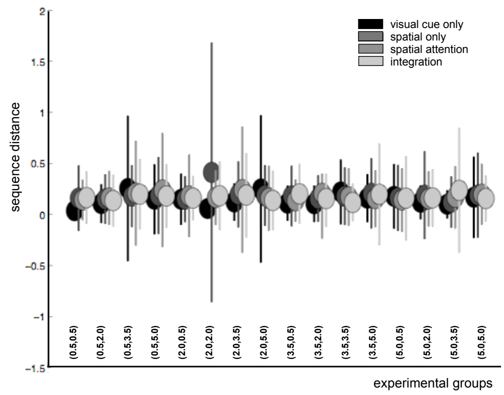

(a)

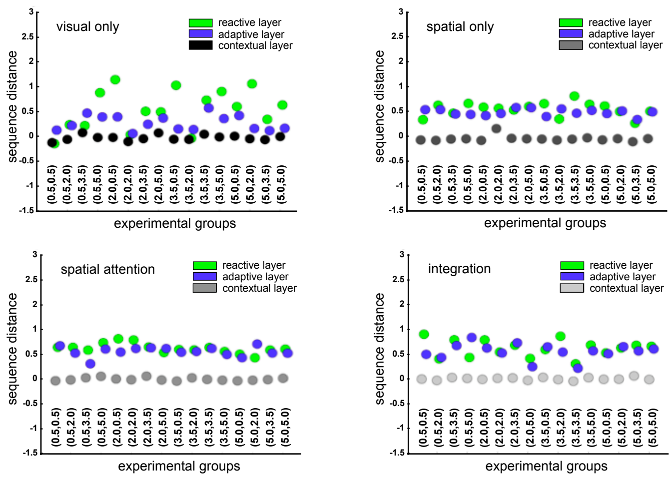

(b)

Fig. 4. Plots represent mean and std of the KL divergence between observation and "optimal" state sequences. a) Optimal state sequences are generated by the HMM model of the contextual layer. From black to light gray: visual cue only, spatial group only, spatial attention and integration group. b) Cross-model KL divergence. Optimal state sequences generated by the HMM model of the reactive(green) and adaptive layer(blue).

of the system is not achieved merely by the activity of its lower layers, namely the reactive and the adaptive one. We analyzed the memory content of each additional layer in the DAC hierarchy and quantified the correlation of the path of the robot with the LTM memory sequences. With the HMM model created by the contextual layer, each fragment from the observed robot path correlates strongly with its "optimal" state sequence. For the HMM models of the lower layers, the correlation is significantly lower as confirmed by anova1 test, $p>0.001$. Here we present as an equivalent measure the KL divergence between path fragments and state sequences, see the Figure 4.

C. Comparison to MC-POMPD

To evaluate our approach in a broader context, we compare the results of the best performing group with the performance of the two lower DAC layers and with the MC-POMPD method [23], one of the known methods from robotics co-munity for learning to act optimally in a partially observable dynamic environment. This algorithm works over continuous action and state spaces and it employs the value iteration, a reinforcement learning algorithm, for action selection. For the purpose of comparison with the DAC, a discrete set of actions (NE, N, NW, W, SW, S, SE, E) is used and distance to the next goal is provided directly as a reward. Each so-called "belief" state, is represented by a set of samples. The actual implementation of how a set of samples is drawn from the state space follows the seminal paper of [24]. Up to K(= 10) similar belief states are stored in a memory. Each memorized belief state is accompanied by the expected reward (utility) for each action. In the MC-POMD method, the belief propagation uses the Monte Carlo sampling. Proprioceptive information and vision are used as input. For a given belief state, an outcome (utility) of all actions is simulated and the action that maximizes the expected reward is selected. A utility value of each action is approximated by interpolation from corresponding utility values of the K nearest-neighbor belief states. The sampling-based Bellman backup is used for their update. The results show that the contextual layer achieves up to three times better performance than the lower layers of DAC hierarchy, and its performance is only ten percent lower then that of the MC-POMDP method.

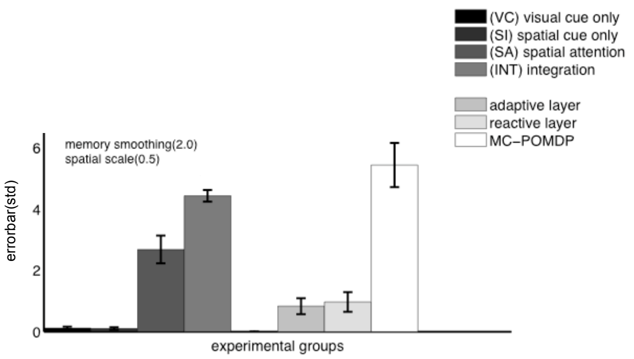

Fig. 5. Comparison of the performance of the reactive layer, adaptive layer, an the MC-POMDP algorithm.

V. Conclusion

Our results show that the approaches which use both the visual and spatial information perform better and, at the same time, they improve the robot's exploratory behavior which we confirmed by inspecting contents of their respective LTM memories. Comparison of the memory content and the performance at the different memory smoothing and spatial scales further suggests that the memory content of the LTM sequences reflects the method by which they are formed.

Both factors, scale of the memory smoothing and the spatial information, have an impact on the performance. A narrow spatial window helps selecting the correct sequences and thus reduces the ambiguity of the information stored in the memory. Widening of the spatial window degrades the performance by introducing additional noise into the system leading to the inclusion of incorrect sequences in the decision making process. For a wide spatial signal, the recall behavior matches the one visual cue only control group. Sequence fidelity, represented by the memory smoothing scale, can be used to counterbalance the performance decrease resulting from increased uncertainty in the spatial information.

If the uncertainty of the spatial information and the memory smoothing are in the lower range, the "integration" approach outperforms the "spatial attention". However, a low variance in memory content and performance across the upper range of memory smoothing and spatial scale of the "spatial attention" suggests the feasibility of a combined approach.

We compared DAC performance with the performance of the MC-POMDP algorithm. The results show that DAC's performance is only ten percent lower then that of the MC-POMDP method. DAC's peculiar way of sampling the input state and of matching learned perception/action pairs could be improved by introducing an additional sequence bias that will directly reflect goal fidelity. However, this result is remarkable considering that DAC does not require a priori global information, i.e. it uses purely egocentric actions vs. allocentric ones used by the benchmark POMDP.

References

[1] Bessiere, P. and the BIBA-INRIA research group. (2003) Survey: Probabilistic Methodology and Techniques for Artefact Conception and Development. INRIA Tech. Report RR-4730

[2] Verma, D. and Rao R. (2006) Goal-based imitation as probabilistic inference over graphical models. (to appear in Advances in NIPS 18)

[3] Rosenblatt, J.K. (2000) Optimal Selection of Uncertain Actions by Maximizing Expected Utility. Autonomous Robots, Vol. 9, pp. 17-25

[4] Voegtlin, T., and Verschure, P.F.M.J. (1999) What can robots tell us about brains? A synthetic approach towards the study of learning and problem solving. Rev. Neurosci. 10(3-4):291-310.

[5] Verschure, P.F.M.J., Voegtlin, T. and Douglas, R.J. (2003) Environmentally mediated synergy between perception and behavior in mobile robots. Nature, 425: 620-624.

[6] Verschure, P.F.M.J. and Althaus, P. (2003) A real-world rational agent: Unifying old and new AI. Cognitive Science: 27, 561-590

[7] Stankovic, M. and Verschure, P.F.M.J. (2007). Optimal multi-sensor fusion in the context of goal-oriented behavior: A robot study, Kuenstliche Intelligenz.

[8] Baker, C. L., Tenenbaum, J. B. and Saxe, R. R. Bayesian models of human action understanding. (2005) Advances in Neural Information Processing Systems 18. (in press)

[9] Griffiths, T. L. and Tenenbaum, J. B. Optimal predictions in everyday cognition. (2005) Psychological Science (in press)

[10] Platt, M. L., and Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. Nature, 400, 233-238.

[11] Massaro, D. W. (1997). Perceiving talking faces: From speech perception to a behavioral principle. Cambridge, MA: MIT Press.

[12] Massaro, D.W., and Friedman, D. (1990). Models of integration given multiple sources of information. Psychological Review, 97, 225-252.

[13] Gallistel, C. R. (1990). The organization of learning. Cambridge, MA: MIT Press.

[14] Roberts, W. A. (1992). Foraging by rats on a radial maze: Learning, memory, and decision rules. In I. Gormezano and E. A. Wassermann (Eds.), Learning and memory: The behavioral and biological substrates (pp. 7-23). Hillsdale, NJ: Lawrence Erlbaum Associates.

[15] Herrnstein, R. J. (1970). On the law of effect. Journal of the Experimental Analysis of Behaviour, 13, 243-266.

[16] Krebs, J. R., Kacelnik, A., and Taylor, P. (1978). Test of optimal sampling by foraging great tits. Nature, 275, 27-31.

[17] O'Keefe, J., and Nadel, L. (1979) Behavioral and Brain Sciences, vol 2(4). Claredon Press, Oxford.

[18] Wyss, R and Verschure, P.F.M.J. (In Press) Bounded Invariance and the formation of place fields. Proceedings of Neural Information Processing Systems (NIPS)

[19] Pearl, J. (1988) Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann, San Mateo.

[20] Rabiner, L.R. (1898). A tutorial on Hidden Markov Models and selected applications in speech recognition. Proc, IEEE, 77:257-286.

[21] MacKay, D.J.C. (2003). Information Theory, Inference, and Learning Algorithms. Cambridge University Press.

[22] Bilmes, J. (1997). A Gentle Tutorial of the EM Algorithm and its Application to Parameter Estimation for Gaussian Mixture and Hidden Markov Models. University of Berkeley, ICSI-TR-97-021.

[23] Thrun S. (2000). Monte Carlo POMDPs. MIT Press, Advances in Neural Information Processing Systems, 12:1064-1070.

[24] Isard M., Blake A. (1998). CONDENSATION - conditional density propagation for visual tracking. Int. J. Computer Vision.