A Novel Framework for Learning Attention Control in a multi-dimensional sensory space

Maryam S. Mirian[1], Majid Nili Ahmadabadi[1,2,3], Hadi Firouzi[1], Babak N. Araabi[1,2], Ronald R. Siegwart[3]
[1]Control and Intelligent Processing Centre of Excellence,
Dept. of Electrical and Computer Eng, University of Tehran
[2] School of Cognitive Science, IPM, Tehran, Iran
[3] ASL, ETHZ, Switzerland

Attention control is a biologically-inspired solution to solve the information bottleneck problem and to make a manageable input sensory space out of a rather distracting one. In fact, attention control explicitly facilitates the main decision making task of an agent. This can be either a learned or a hard-coded control strategy. However, if the main decision making task is difficult, the latter is not a consistent solution. This is the reason we try to learn how to control the attention of an agent. But, using attention control learning brings forth a new optimization problem. Now, we are confronted by two tightly coupled learning problems. We transfer the space in which we want to solve learning problems in order to decouple thus facilitating the solution.

In brief, a novel mechanism to transfer the learning of attention control to action-value space while also learning a challenging decision making task is proposed. A three-layered framework for learning attention control in action-value space (decision space) is presented. The originality of this transfer is that the agent gradually learns to control its attention in parallel with learning a demanding decision making task (coarse driving) in the decision space. We model the learning to drive as a task which needs controlled eye movements on a wide front window. In other words, one can assume that an attentive driver is one who can professionally (as quickly as possible) focus his or her eyes on the most important object or event at the appropriate time.

The driver robot's mind is designed as a cooperative Multi-agent system. It consists of several local decision experts (LDE) and a final decision maker (FDM). Each LDE is an expert on one spatial area of a distracting perceptual space. There is an overlap among the spaces that LDEs are trained over. FDM does the final motor action. Training LDEs is done imitatively in Layer 0. This training continues in Layer 1 using an explicit expertness indicator mechanism to fuse the decisions of LDEs. Finally, learning how to fuse decisions of LDEs to find the final optimal motor action is done in Layer 2. This final layer is responsible for learning attention control. Since the state space is continuous, a Bayesian continuous Reinforcement Learning approach is utilized in all three mentioned layers.

- Layer 0: Initial Imitative Training of LDEs: In this phase, a teacher drives in front of LDE's eyes. So, they observe both the environmental state and the corresponding action that the teacher does and the association between them is reinforced.
- Layer 1: Training of LDEs using a fixed expertness indicator measure for fusion: All LDEs observe the environment in a parallel manner. They all propose their ideas about best motor action to be performed. A Fusion mechanism indicating maximum expertness is used to find the final motor action. When it is decided and performed, the reinforcement signal is taken and all LDEs update their own learning weights. Figure 1 shows a schematic view of this layer, assuming that we have either three different modalities or some spatial areas inside a modality: S_1, S_2, S_3.
- Layer 2: Attention Control Learning by consulting LDEs: Let's start with the traditional form of attention control learning question: "if we have at most N sensors to perceive the environment, which n ones are more informative yet most cost effective to be turned on?" In the new space, we have M LDEs (replacing N original sensors, $M \leq 2^N-1$) So, at each step of driving task, FDM tries to answer: Which m (out of M) LDEs to consult with (each consultation imposes a processing or time cost) in order to find a more descriptive (i.e. less ambiguous) state with the lowest accumulative cost? Figure 2 roughly demonstrates what happens in this layer.

Learning attention control in decision space has some interesting advantages over learning attention control in perceptual space. The major ones are sharing the common decision space among LDEs, utilizing not necessarily similar learning algorithms for LDEs and finally making a more confident decision by FDM.

It has been shown that the attention control is learned as well as the main robot's task (autonomous driving) while the accumulative cost is kept minimized. The real miniature vehicle is an e-puck robot which drives smoothly in a round road while avoiding stationary and mobile obstacles.

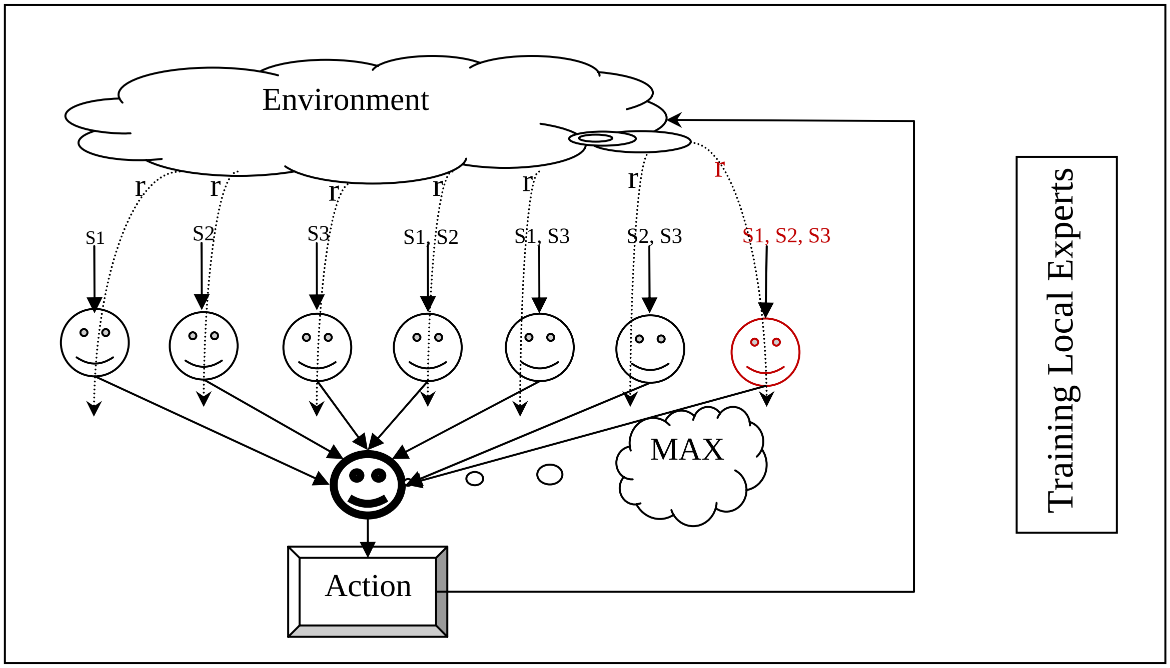

Figure 1- Layer 1: Training of LDEs using a fixed expertness indicator measure for fusion

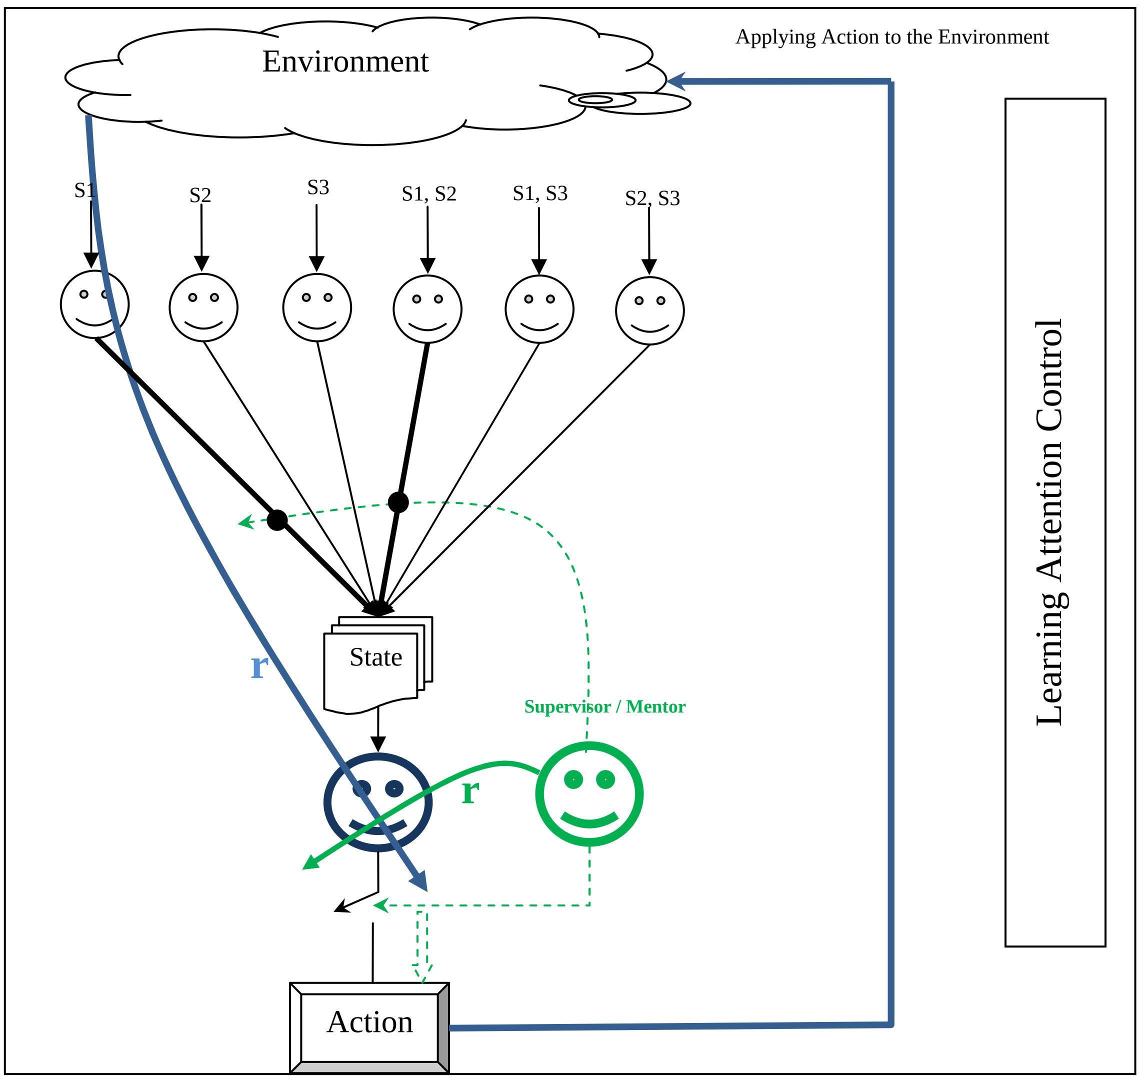

Figure 2- Layer 2: Attention Control Learning by consulting LDEs