Pro-active agents with recurrent neural controllers

Guido de Croon
IKAT,
Universiteit Maastricht

To obtain my M.Sc. in Knowledge Engineering at the IKAT in Maastricht, I have performed research on Evolutionary Robotics, a methodology that optimises robot controllers with evolutionary algorithms. Part of the research for the thesis has taken place at the CNR (Consiglio Nazionale della Ricerca, Rome), enabling me to collaborate with Dr. S. Nolfi, an expert in the area of Evolutionary Robotics. My supervisors at Universiteit Maastricht were Prof. dr. H.J. van den Herik and Prof. dr. E.O. Postma.

In my thesis I have focused on a bottom-up approach to artificial intelligence, called Embodied Cognitive Science. As a consequence of the bottom-up approach, research has emphasised reactive agents. These agents always respond in the same way to the same sensory inputs. It has been proven that these simple agents can perform complex tasks by exploiting sensory-motor coordination: the agents use their actions to obtain advantageous sensory inputs. To reach higher agent capabilities, research focus is currently shifting from reactive to pro-active agents. The actions of pro-active agents do not only depend on the inputs, but also on the "internal state". An internal state represents a form of memory of the agent's past sensory inputs. Typically, a pro-active agent has a recurrent neural controller. A large variety of such controllers has been proposed, but it is not yet clear in what way the different neural controllers influence the agent capabilities. Therefore, it is also not clear what the next step should be in the bottom-up approach to artificial intelligence.

Our research question was: how are a pro-active agent's capabilities influenced by its recurrent neural controller? To answer the research question, five typical recurrent neural networks were applied as controllers of a simulated Kephera robot in three different robotic tasks.

Our conclusion is that the capabilities of a pro-active agent are determined by the mechanism that realises an internal state in its recurrent neural controller. We discern three such mechanisms: recurrency, neural inertia, and adaptable time delays on the neural connections. Neural controllers that employ the mechanisms of neural inertia or adaptable time delays, lead to agents that can exploit regularities on variable time scales. Specifically, these mechanisms offer the agent the ability to determine when sensory inputs influence the outputs. The mechanism of recurrency alone is not sufficient to obtain this ability.

Illustrative example

With an example we illustrate why an inability to determine when sensory inputs influence the outputs leads to lower agent capabilities. One of the tasks requires from the agents that they self-localise in an environment at various speeds. The figure shows two screenshots of the environment. The left part of the figure involves an agent controlled by a Nonlinear AutoRegressive model with eXogenous inputs (NARX), while the right part involves an agent controlled by a Continuous Time Recurrent Neural Network (CTRNN). The agents have to indicate with an output neuron whether they are located in the light grey room or the dark grey room. To exemplify the limited capabilities of an agent that cannot determine when the inputs affect the outputs, we discuss the way in which the agents notice that they enter the bottom room.

The CTRNN-agent (right) uses its internal state to indicate in what room it is located. In particular, its fourth hidden neuron is part of its internal state and determines the activation of the self-localisation output. The CTRNN-agent uses its neural inertia to notice when it enters a new room. Neural inertia is a mechanism in which the neural activation changes with a certain speed, determined by evolution. We discuss the transition to the bottom room. The activation of the agent's fourth hidden neuron slowly decreases in a corridor. Since the activation increases again in a turn, the agent can be said to "measure" how long the current corridor is. The agent can determine when it enters the bottom room, since there is only one corridor in the environment that is long enough to allow the activation of the hidden neuron to decrease to zero. In particular, the self-localisation output indicates the bottom room when the fourth hidden neuron decreases below 0.5. The fourth hidden neuron is shown in the bottom of the figure and a dashed circle shows the moment in which the agent indicates the bottom room.

The NARX-agent (left) uses its self-localisation output, which is part of its internal state, to "remember" in what room it is. However, the agent does not use its internal state to recognise when it has to change its localisation. Instead, it uses sensory-motor coordination to obtain unambiguous information about its location. We again discuss the transition to the bottom room. In the corridors the NARX-agent makes heavier "swings" than the

CTRNN-agent. The agent's movements in the long corridor illustrate this best. The swings of the NARX-agent serve to determine approximately where the bottom room is. Namely, there is only one place in which the agent approaches the wall to its left without being in a turn: in the middle of the long corridor. In that moment the internal state changes abruptly, as illustrated by the change in activation of the self-localisation neuron in the bottom of the figure.

The experiments suggest that NARX-agents have lower capabilities than CTRNN-agents, because no successful NARX-agents have been evolved at higher speeds. The higher required speed restricts the possibilities to make swings in the corridors and NARX-agents do not produce a strategy with slowly changing neural activations. The CTRNN-agent however can still perform well at the task at a higher speed, since its neural inertia is adjusted to decrease faster. In other words, the agent can determine the time scale on which the inputs experienced in the corridor lead to an indication of the bottom room.

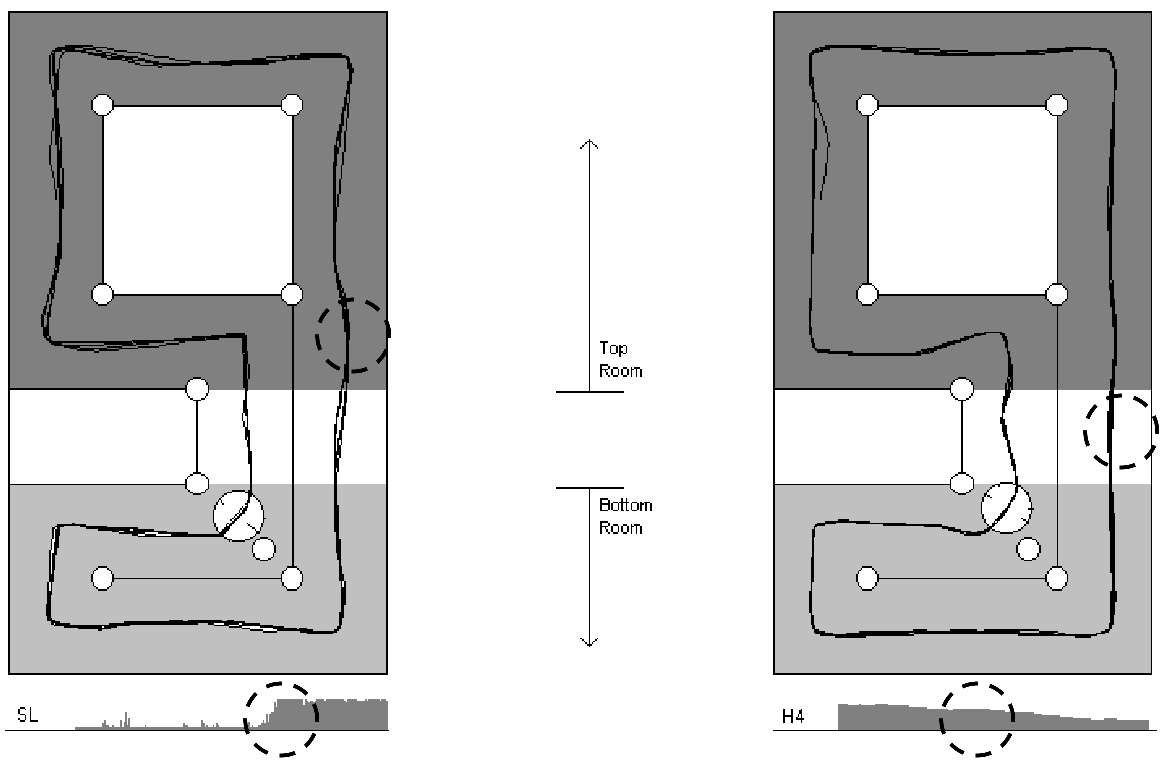

Figure: The left agent has a NARX-network, the right agent a CTRNN. Lines represent walls, small cylinders are obstacles. The dotted circles indicate when the agents change their internal state. Below the environments, the most important neural activations are shown. 'SL' stands for the self-localisation neuron of the NARX-controller, 'H4' is the fourth hidden neuron of the CTRNN-controller.