Old Actions in Novel Contexts — a Cognitive Architecture for Safe Explorative Action Selection

Malte Schilling[1]

Abstract. When running into a problem for the first time a behavior control system is supposed to come up with a sensible response. Most importantly the selected action should be safe and not dangerous for the system. While adaptive behaviors can deal with broad variations of contexts, for truly novel situations the system has to explore which action might be suitable. In a cognitive system this exploration can be realized in a safe way through internal simulation. Internal simulation allows the exploration of the consequences of possible actions without actually taking the risks. Here, a control approach for six-legged walking is presented which is biologically inspired by detailed work on the walking of stick insects. The underlying biomimetic control architecture has been extended towards a cognitive architecture that allows for planning ahead. In order to exploit these predictive capabilities in novel situations, the behavior selection part of the system has to be extended. The goal is to allow for trial-and-error testing of behaviors even though the current sensed context is not encouraging the use of that particular behavior. This article introduces a three layered neural network approach that realizes explorative action selection. First, this neural network selects step-by-step single behaviors, exploring the space of possible actions from a given starting point. Second, it remembers which behaviors have been chosen already and should not be selected again. This network is part of a cognitive expansion which allows the hexapod robot to overcome novel problems during walking through planning ahead.

1 Introduction

The main tenet of embodied cognition is the idea that higher cognitive capabilities are rooted in the sensorimotor functions and cannot be understood separated from those [2, 3]. Following this approach, the presented work aims at a minimal cognitive system which is constructed in a bottom-up fashion. First—and on the lowest level—this system deals with motor control for a hexapod walker. The organization of the action selection scheme is biologically inspired and is characterized by its decentralized nature [17] . Second, the control structures of the behaviors as such which are in this case realized as neural networks.

The central question of the presented research is how can such a simple behavior-based system become a cognitive one and what kind of mechanisms are required? One part of the answer involves the notion of internal models which is crucial for higher level cognition. But, from my perspective it is not the existence of internal models which make a system a cognitive one [6]. Instead crucial for a cognitive system is when existing internal models can be exploited in planning ahead. One example for the case of a six-legged walking controller is the control of the stance movement which involves multiple legs and their joints. This requires the coordination of all the single joint movements which makes it a complex problem and requires explicit coordination. An internal body model provides this form of explicit coordination and such a model has been introduced into the system. The internal body model is realized as a recurrent neural network and has shown to be quite flexible and predictive.

For the minimal cognitive system approach the system has been therefore extended in a way that exploits the underlying internal body model [20]. As a main idea, the predictive capabilities of the internal model are utilized in a decoupled mode [11, 24] to provide a form of internal simulation which realizes planning ahead. Here, I follow McFarland and Bösser [12] who define a cognitive system as one that is able to plan ahead. In the presented case, cognition is understood as the ability to plan ahead [12] by a means of an internal simulation [11] relying on an internal representation [23, 10]—starting with a model describing the spatial and dynamic relations [1] of the own body [7]—which is grounded in embodied experiences [9]. As a consequence, it did not become necessary to introduce quite complex internal models into the system in order to become cognitive. Instead, crucial is the simple mechanism of decoupling and a trial-and-error approach to action selection.

This article addresses the question how the selection of actions can be broadened and not be restricted to the usual (programmed or learnt) sensory contexts. I will explain how the exploration of the action space can be realized through a simple three layered neural network structure. For details on the subsequent safe internal simulation of selected behaviors see [20]. The paper is structured in the following way. The second section summarizes the behavior-based approach and gives an overview of the cognitive expansion. In particular, the three layers required for explorative action selection are explained in detail. The third section shows results for the activation of those three layers highlighting how these layers realize action selection. Last, the discussion will connect with related work and conclude on future work.

2 Control of Hexapod Walking

The control system for the hexapod walker is biologically inspired by detailed work on the walking of stick insects [17]. There are two general problems which have to be dealt with. First, the selection of the appropriate movement for a given context. Second, the execution of the selected movements. Here, I will only deal with the first problem and will not address the details on how the specific movements might be carried out. In case of the stick insect there are mainly

[1] Center of Excellence 'Cognitive Interaction Technology' (CITEC), Bielefeld University, P.O. Box 10 01 31, D-33501 Bielefeld, Germany Email: mschilli@techfak.uni-bielefeld.de

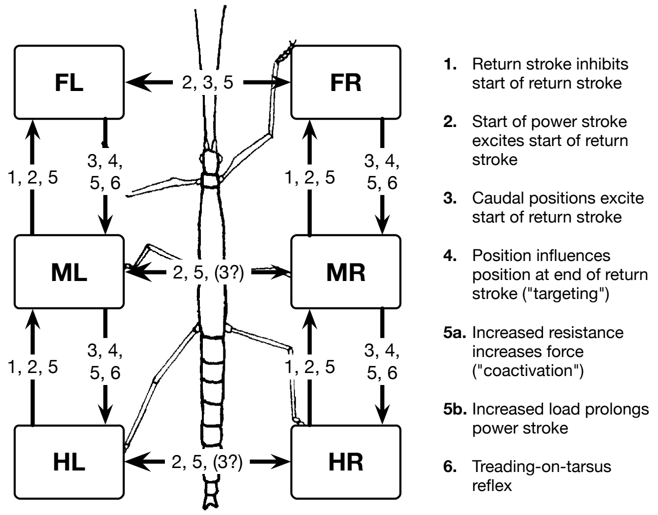

Figure 1. Decentralized control architecture. Each leg is controlled by an individual controller (shown as a box) which selects the appropriate behavior for the given context. The coordination influences are obtained from detailed experiments on walking in stick insects. For details see [17].

two basic movements concerning a single leg. On the one hand, the stance movement and on the other hand, the swing movement. Details for the swing movement can be found in [22]. During the stance movement multiple legs are connected through the ground in closed kinematic chain making this a difficult control problem. It becomes necessary to coordinate the movement of at least nine joints. This requires some form of explicit coordination. A recurrent neural network has been introduced as an internal body model which is mediating the movements of all participating joints [19, 14] . Importantly, while the model is solving the inverse kinematic task during the stance movement, it is a flexible and predictive model as well [15]. Therefore, it can be used for testing behaviors in internal simulation. Following Steels [23], the body model is grounded as it has co-evolved in service of a particular behavior. As a next step it can now be exploited for planning ahead by means of internal simulation [20].

In the following, I will first introduce the underlying biologically-inspired control system which selects the appropriate motor primitive depending on the given context. Second, I will extend this system towards a cognitive system which allows the testing of existing movements out of their original context which means that these movements can be selected independently of the sensed state. As this selection might potentially lead to a dangerous situation, the cognitive part of the system is not applying chosen movements directly on the real robot. Instead, an internal simulation is used as an estimate of the consequences of the application of that movement which allows the system to make an informed decision about the movement chosen. Last, the main and novel focus of this article lies in the simple neural network structure that is driving the selection process of the movements independent of the given sensory context.

2.1 Behavior-based walking controller

One main characteristic of the biologically inspired controller is its decentralized nature. There is a single controller for each leg that selects which action should be carried out. These controllers are only loosely coupled through local coordination influences (Fig. 1). The overall walking behavior emerges as a result from the individual controllers [17].

The controller on the leg level is realized as a neural network consisting of so called Motivation Units (Fig. 2, [18]). Motivation Units are simple neuronal units which sum up the incoming inputs and have a linear characteristic. As an output function, a Motivation Unit restricts its activation value to be in the range between zero and one. A Motivation Unit can be further associated with a movement which should be carried out. In that case, the Motivation Unit represents if the specific behavior is activated and to what extent. The connected motor primitive allows the control of the motors of the walking system. In this way, a Motivation Unit is representing an activation or motivation of the specific motor program. Motivation Units can be further organized into heterarchies. In the case of the hexapod walker, there are mainly two different movements: the stance and the swing movement (see Fig. 2, orange units in central part of the figure). Each of the respective behaviors is represented by one Motivation Unit. As only one of those behaviors should be active at the same time, those two behaviors are mutually coupled through inhibitory neural connections. As a consequence, when one of the neurons is active it suppresses the other one and in consequence only one of the two behaviors can be active at any given time. The walking system has been extended to allow for forward as well as backward walking. As a consequence, there are two swing and stance Motivation Units, one for both directions.

In order to further stabilize the ongoing behavior, there are recurrent connections (not shown) enforcing the currently active behavior. Therefore, sensory events are required to switch between behaviors and these define the contexts of the motor behaviors. On the one hand, while a leg is in swing mode it is lifted off the ground and moving to the front of the working range (in the case of forward walking). When moved to the anterior part of the working range the movement produced leads the leg downwards until it touches the ground. The sensed touchdown enforces the stance movement and in this way invokes a switch to the stance movement. On the other hand, in forward walking during the stance movement the leg is supporting the body and moves backward until the leg reaches its posterior extreme position (PEP). When the leg reaches the PEP the stance motivation is inhibited and the swing movement is activated (Fig. 2, Δ-PEP Motivation Unit). The variable PEP can be influenced by the neighboring legs through the coordination rules. For example, the first rule simply states that while a posterior leg is in swing mode and therefore in the air, for the anterior leg the PEP is shifted backwards which prolongs the stance movement of that leg providing stability until the other leg touches down (Fig. 1).

This control system has been termed Walknet and is a behavior-based control system which has been demonstrated multiple times in dynamic simulation [16, 4] as well as on different robots [13, 21] to produce adaptive stable walking in all kind of challenging scenarios.

2.2 Overview Cognitive Expansion

The Walknet control approach is an embodied approach and the activation of behaviors is driven by sensory influences. As a consequence, the sensory states define contexts in which behaviors should and could be activated. The Motivation Unit network provides an extendable architecture to further incorporate additional behaviors (as already done for the case of backward walking). Importantly, usually when designing such a system one is interested in, first, that only one behavior is active for each leg at any given time. Second, that a behavior becomes only active in a quite well-defined situation for

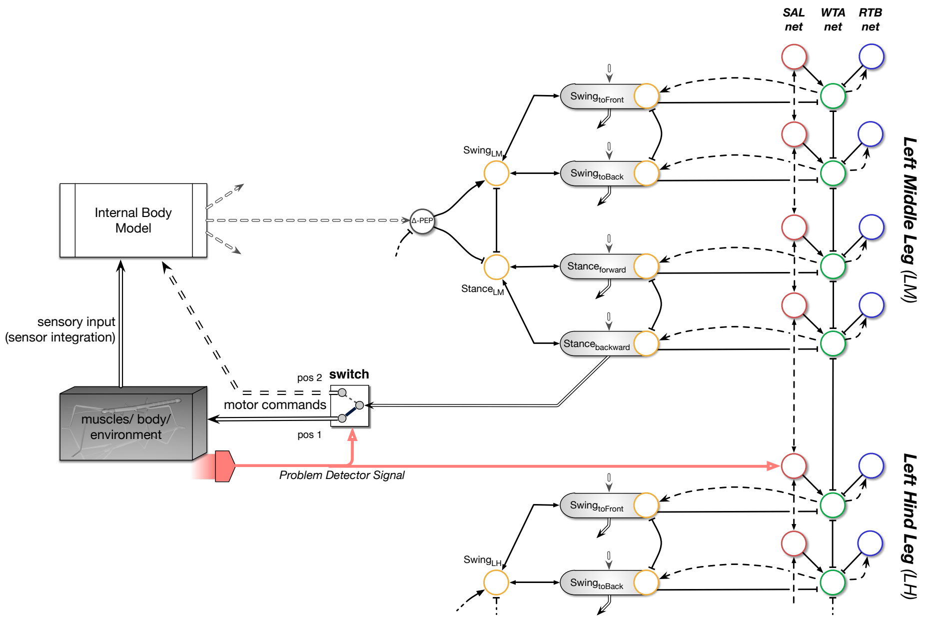

Figure 2. A single leg controller (left middle leg) is shown. The main control structure for the left middle leg is given (center, upper part). There is each a stance and swing movement for forward and backward walking. Stance and swing Motivation Units are connected through mutual inhibitory connections. On the right, the cognitive expansion is shown which allows for the selection of a behavior out of context. On the left, the body and the internal model are shown. In normal walking the motor primitives send signals directly to the body (switch is in position one). When a problem is detected (see red problem detector), the current behavior is stopped and the search for an alternative behavior is started in the cognitive expansion on the right. The problem detector induces activity in the cognitive expansion structure as it also provides information where to start the search. Secondly, the problem detector flips the switch to position two and as a consequence the motor commands are not routed to the body, but instead are routed to the body model which provides predicted consequences of the selected motor program. This predictive internal simulation can be utilized to decide if the selected behavior causes any further problems or could be safely selected for application on the real robot. For details see [20].

which it is appropriate. This has allowed Walknet to deal with quite diverse and challenging scenarios. But, when running into a novel situation this approach might be too restrictive and no behavior might become active.

As one such example it is possible to disturb walking by adjusting the leg positions during the swing movement (for details see [20]). When a right hind leg and the contralateral middle leg are moved far to the front, the whole system can become unstable as soon as the left hind leg starts a swing movement. The system might topple over (Fig. 3, in contrast to the robot, this case does not pose a threat for the insects as the insects are using their tarsi to hold on to the ground which prevents toppling over). In the robot case, where there are no adhesive structures at the leg tips, a problem detector has been implemented which recognizes such instabilities and stops walking any further. In the behavior-based approach the activated behaviors run the system into an impasse. The system gets stuck.

But a solution to such a case might be quite simple, for example through repositioning a single leg. Therefore, the behavior-based system has been extended in such a way that when the system runs

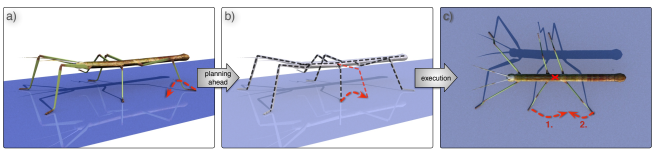

Figure 3. Awkward posture: when the middle left and right hind leg are moved further to the front, the hind left leg cannot be lifted anymore without the system toppling over (see subfigure c) where the x is marking the center of mass).

into an impasse it loosens it restrictions on which behavior to select. But, importantly, simply choosing any behavior might be dangerous. It is therefore not a good idea to just execute any randomly selected behavior. Therefore, the control system is extended towards a cognitive system as it allows for planning ahead. As part of the cognitive expansion a behavior is selected (possibly out of context), but it is not immediately applied to the real robot. Instead, the predictive capabilities of the walking system are exploited and for that particular selected behavior the continuation of the overall movement is predicted over time in internal simulation. The internal simulation provides a prediction of what would happen if that particular behavior would be selected and makes an informed decision on whether the system would become unstable or not.

For the chosen example, one possible solution would be to reposition the middle left or hind right leg in order to provide support when lifting the hind left leg. Details for this example case have been described in detail in [20]. Fig. 2 provides an overview of the overall system and how the cognitive extension is build on top of the behavior-based system. In brief, the problem detector, first, flips the switch. As a consequence, motor control signals are now rerouted towards the internal body model which is run in a loop and provides a prediction of what would happen if that particular behavior would be chosen. As the internal body model reflects the real robot, it also contains a problem detector (not shown in the figure). Therefore, instabilities in the internal simulation can be detected, too. Second, in internal simulation different motor programs should be tested one after another until a solution is found. Importantly, the selection of the motor programs should not be restricted to the original intended context of these behaviors. This article focusses on how this selection is performed through the three additional layers shown on the right in

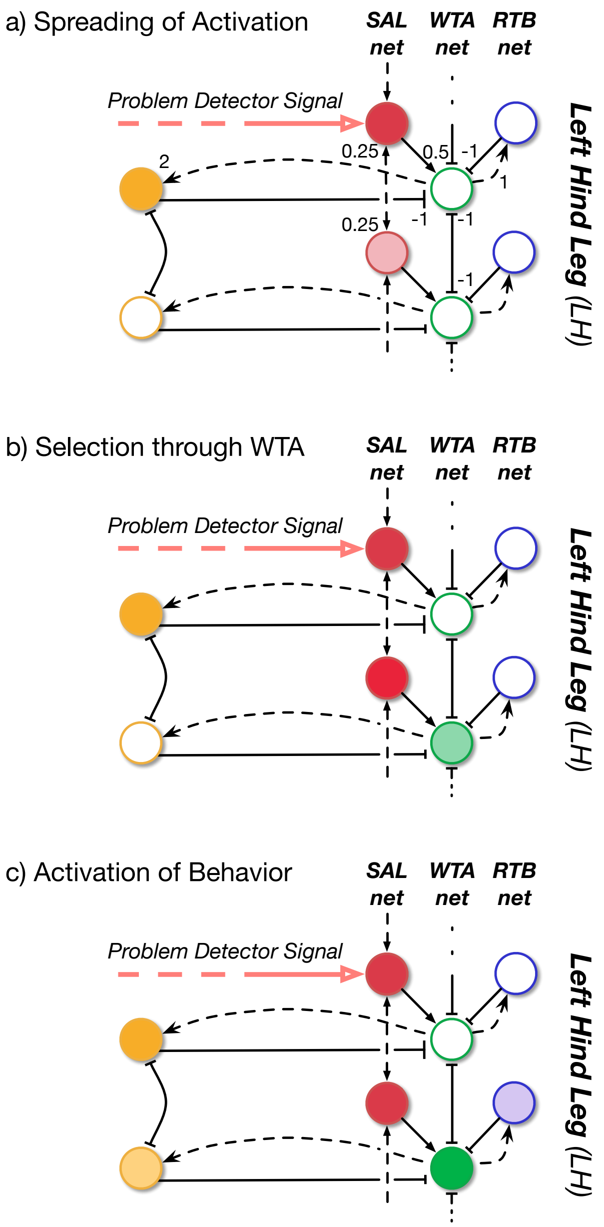

Figure 4. The three layers during different processing phases are shown. a) shows how activation (strength of activation is given as a filled unit) is induced in the SAL layer and spreads through neighboring units (weights of the connections are shown). In b) the selection of a single unit inside the WTA layer is shown. This selection, on the one hand, activates the corresponding Motivation Unit on the left and, on the other hand, the RTB unit on the right which stores that the unit became already active (c). Dashed connections are modulated, i.e. only active during the corresponding phase.

Fig. 2. For further details on the internal simulation, the timing of the overall system during the different processing stages as well as the problem detectors see [20].

2.3 Explorative Action Selection

The cognitive expansion introduces three additional layers of Motivation Units. For each Motivation Unit representing a behavior one corresponding unit is introduced in each of the three layers (Fig. 4). The overall task of the action selection part of the cognitive expansion is to select one single behavior at a time which should be tested in internal simulation. First, there is the spreading of activation layer (SAL net) which drives the exploration. Second, the winner-take-all network (WTA net) guarantees that only a single unit becomes active. The third layer (remember-tested-behavior or RTB net) keeps track of which behaviors already became active and should not be selected a second time.

The exploration of the available action space requires these three additional layers and the explorative selection is accomplished in two phases (the timing of the these phases is addressed in [20]). When the system runs into a problem this is detected by a problem detector, e.g. the system becomes unstable when trying to start a swing movement with the left hind leg. This information on the cause of the problem is exploited as a heuristic where to start the exploration. In the SAL layer (Fig. 4, a) an activation is induced in the Motivation Unit connected to the Motivation Unit of the swing movement of the left hind leg. This is used as a starting point for the exploration which is realized as a simple spreading of activation. In the simple approach presented, the spreading of activation follows the structure of the controller and the animal. All Motivation Units in the SAL net are stringed together and each one has two neighbors. There is no hierarchy assumed here, but in future work this could be exploited. Right now the only assumption is that when a problem is caused in the left hind leg movements of legs close by should be considered first. Fig. 4, a) shows how the activation spreads inside the SAL layer during this phase. There is a small stochastic portion in the activation of the units in the SAL layer as noise is added (std. dev. ± 0.10) to the SAL units.

In the second phase (Fig. 4, b), the selection should be narrowed down to a single Motivation Unit as only a single behavior should be tested in the internal simulation later-on. This is realized inside the WTA layer. Each unit gets activated by the corresponding unit in the SAL layer. Inside the WTA net, the units are inhibiting each other. As a consequence from this winner-take-all structure only the activation of one unit will survive over time. The winner-take-all net gets two additional inhibitory influences. On the one hand, it is not necessary to select the already active behaviors as those have caused the problem. On the other hand, the corresponding RTB units are memory units which store when a unit has been selected. For an earlier selected unit the activation of the RTB unit effectively excludes the corresponding unit in the WTA layer from taking part in the competition. When the WTA net has settled and a single activation is left, the winner-take-all stage is completed.

Last, after a single Motivation Unit has been selected the corresponding Motivation Unit representing the behavior as such will be activated and the behavior will in this way be tested in internal simulation (Fig. 4, c), orange unit on the left side becomes active).

3 Results

The previous section has sketched out the general idea of the explorative action selection. This section will provide example runs showing how the simple three layered expansion is sufficient to drive action selection. I will concentrate on the three layers only and use eight units in each of the layers as an example. There is an activation induced in the third unit (counted from the top) which indicates where the problem has occurred and where to start with the spreading of activation.

The results will be presented in three steps. First, the spreading of activation will be explained. Second, for one row of units the activation over time will be shown in detail. At last, the activation and

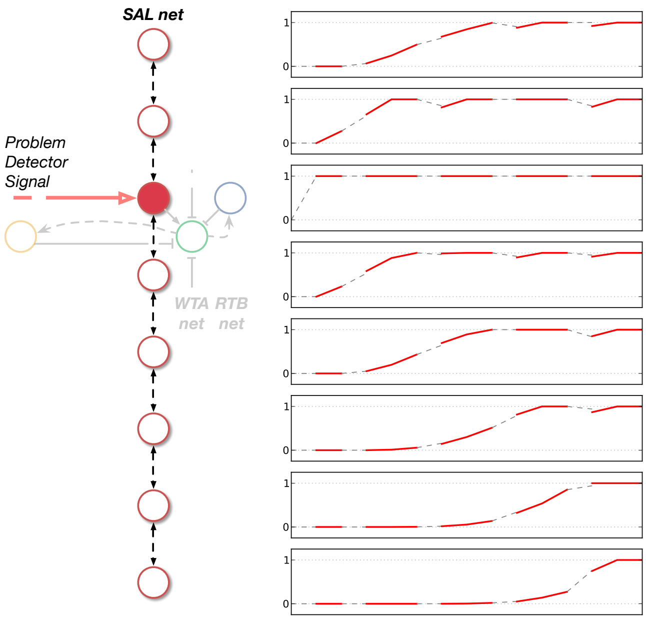

Figure 5. Spreading Activation Layer. Shown are eight units and the connections to the neighboring units. These connections are modulated and only active during the spreading activation phase. An activation is induced in the third unit from the top (input from a problem detector). On the right, the activation of each unit over time is shown. The red parts of the graphs show the spreading of activation phase. The activation spreads from the unit activated first (marked red) to both sides

selection over time will be shown for the complete example.

Fig. 5 shows how the induced activation is spreading from one unit to the neighboring units. On the right hand of the figure the activations are shown in red. The focus is on the spreading phases and not on the time between two of those phases (like the following winner-take-all phase and the internal simulation itself). The other phases are cut out and summarized by the dashed grey lines in between. Importantly, the spreading of activation is modulated only during the spreading phase (for details on the overall timing see [20]).

As a second perspective I want to focus on the activations of three corresponding units spanning the three layers. One such row is shown in Fig. 6. The row corresponds to the second row shown in Fig. 5 which is directly next to the row in which an activation is induced during the spreading of activation. For the different units the activation is highlighted by solid lines showing when these units are active. First, the unit in the SAL layer gets activated by the neighboring unit. In the second phase—the selection phase—the corresponding (green) unit becomes activated and takes part in the winner-take-all competition, but loses during the first run. In the following spreading of activation phase the unit in the SAL layer gets activated even more (the selection and testing of the behavior in internal simulation is left out). As a consequence, the unit wins during the next selection phase. After the selection phase the winning unit also activates the corresponding unit in the RTB layer (this connection is also modulated). In following selection phases this activation of the RTB unit effectively excludes the unit from taking part (through the inhibitory connection).

Last, Fig. 7 shows the whole network structure and how this allows the successive selection of one behavior after another. The different phases are clearly visible. First, activation is spread (shown as the continuous red line). Second, from the active units one is chosen

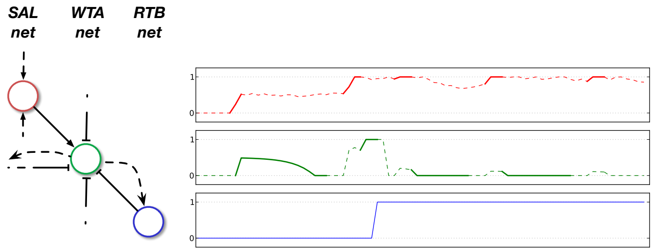

Figure 6. Activation of the three layers for a single row which is connected to a single Motivation Unit representing a specific behavior. Here, the second unit of Fig. 5 is shown. The top panel shows the activation of the SAL unit, the middle panel shows the WTA unit and the bottom panel shows the RTB unit. In the top and middle panels the continuous line parts illustrate when the respective phase is ongoing.

through WTA competition (shown as the continuous green line). Afterwards, the process continues as the activation is spread even wider and the next unit is selected. The figure shows how step-by-step one unit after the other is selected beginning with the units close to where activation was induced. After a unit in the WTA layer has been selected, the corresponding RTB unit becomes active (indicated by the shaded blue area). As this activation is inhibiting the WTA competition, in this way the corresponding unit is excluded from the selection process.

4 Discussion and Conclusion

This article proposed how explorative action selection can be realized in a simple, three layered neural network structure. The neural network structure is part of a cognitive expansion which allows a hexapod walking system to plan ahead in novel and problematic situations [20]. Planning ahead is realized as an internal simulation exploiting the predictive capabilities of an already present internal body model. Initially, the body model is required for the control of the stance movements and coordinates the joint movements of all participating legs. In planning ahead, the internal body model is used as a simulator in internal simulation. This allows the testing of different behaviors and the prediction of consequences of the application of a single behavior which realizes a form of detailed motor planning. When such a cognitive system runs into a novel problem for which there is no suitable behavior that becomes activated by the sensed state, it can safely try to broaden the scope of the already present behaviors and test a behavior out of its original context safely in an internal simulation. The internal simulation provides an assessment if the selected behavior is helpful and in particular non dangerous. Here, I presented how this behavior selection can be driven in an explorative manner using a simple three layered neural network structure. This cognitive expansion starts to search for behaviors which are—in this case morphologically—close by and ensures that only a single one is selected and that each behavior is selected only once.

The idea of internal simulation in robots has been applied already successfully by Bongard and colleagues [5] . The goal of their approach was to evolve a walking behavior for a quadruped robot. The evolutionary approach consisted of two stages. First, through "motor babbling" an internal body model was evolved. Second, the body model was used to evolve in internal simulation a walking behavior. As both stages alternated, this approach allowed the robot, on the one hand, to come up with a successful walking behavior; on the other hand, the system was able to adapt its body model to morphological changes. In contrast to the approach presented here, their approach

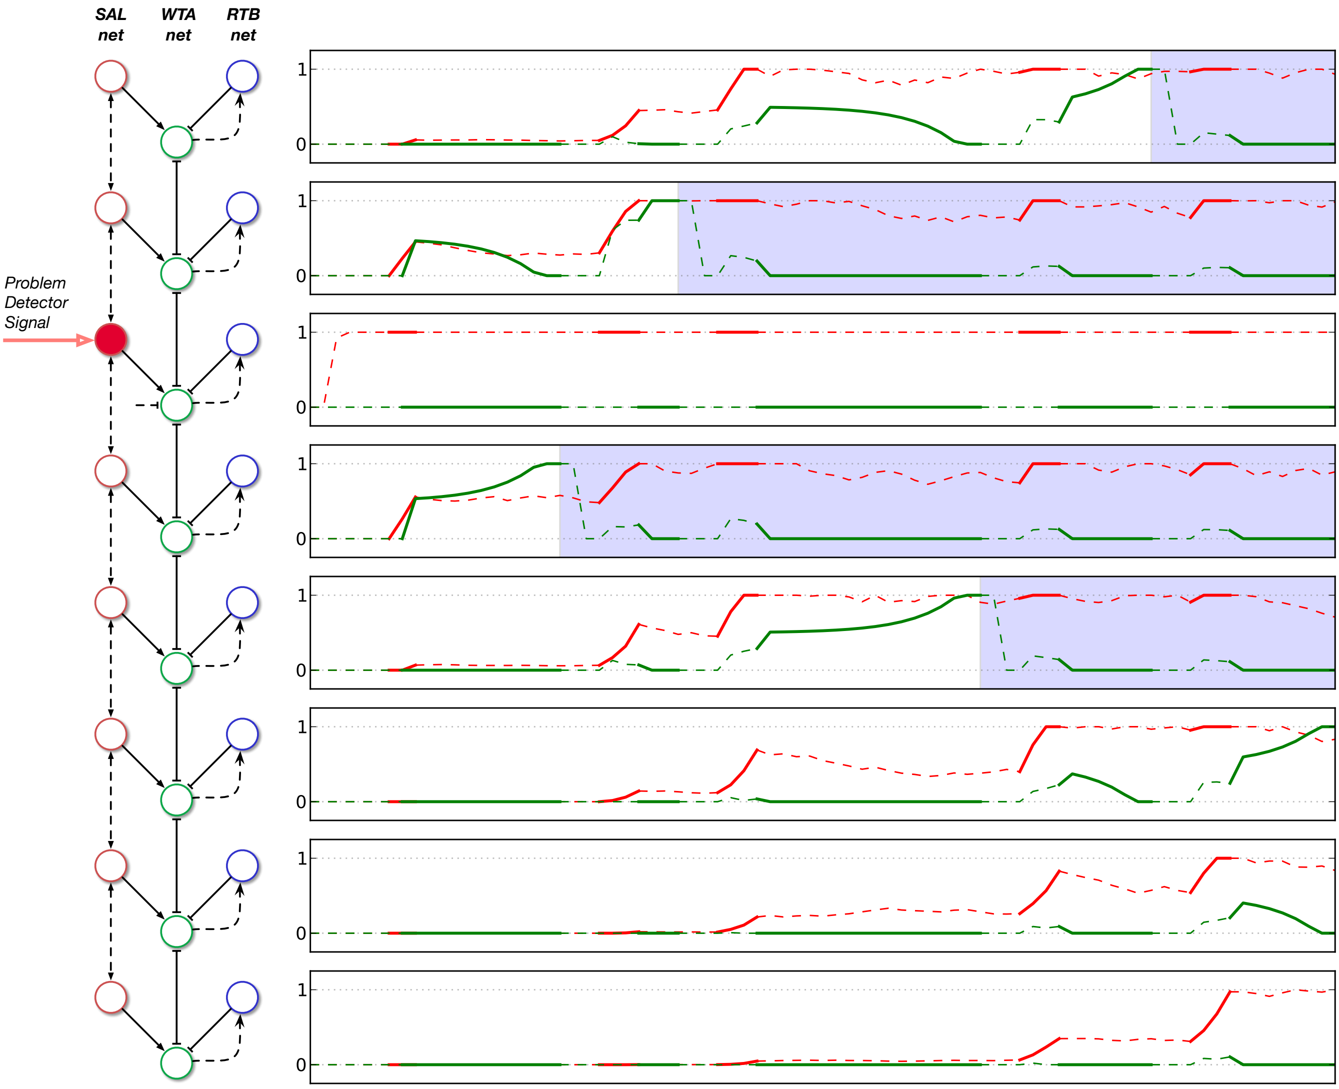

Figure 7. Explorative selection. For each triplet of units the activations are shown on the right. SAL units are drawn in red and WTA units in green. The activation of the RTB unit is marked by the light blue area which indicates that this set of units has already been selected once.

did not start from a biological perspective and did not take into account the decentralized nature of such systems. Second, they only dealt with adapting a single behavior while the presented approach is explicitly dealing with the question of behavior selection.

More recently, Cully et al. [8] presented an exciting approach to locomotion in a hexapod robot. The model also uses a trial-and-error mechanism to test new behaviors. In contrast to the approach presented here, these behaviors are tested on the real robot and not in internal simulation. This becomes possible as they are not focussing on the question of action selection. Instead, their approach is based on a behavioral parametrization which describes the current behavior in a low-dimensional space. During trial-and-error the parameters are adapted. The mapping from the parameters to the very high dimensional space of possible controllers as such is solved through a previous offline optimization which guarantees to lead to non dangerous behaviors. Again, this differs from the reaCog approach as it is not considering the question of behavior selection or taking the organization of action selection into account.

For the future it is one goal to incorporate influences from the mentioned approaches. The current system is dealing with the question of action selection and could be extended in a similar way to also address how existing behaviors could be successfully adapted in specific situations. Another aspect is the organization of the action selection scheme. Currently, the Motivation Units are organized simply on a string just reflecting the basic morphology of the animal. The main idea is that when a problem occurs in a specific leg, it seems reasonable to start testing movements in legs close-by. But this structure could be extended and could incorporate quite diverse heuristics. Third, learning of successful selected behaviors and ideally the possibility to transfer such learned knowledge is another goal.

ACKNOWLEDGEMENTS

This research/work was supported by the Cluster of Excellence Cognitive Interaction Technology 'CITEC' (EXC 277) at Bielefeld University, which is funded by the German Research Foundation (DFG).

REFERENCES

[1] C.A. Acosta-Calderon and H. Hu, 'Robot imitation: body schema and body percept', *Journal of Applied Bionics and Biomechanics*, **2**(3-4), 131–148, (2005).

[2] L.W. Barsalou, 'Grounded cognition', *Annual Review of Psychology*, **59**, 617–645, (2008).

[3] J.R. Binder and R.H. Desai, 'The neurobiology of semantic memory.', *Trends in cognitive sciences*, 1–10, (October 2011).

[4] B. Bläsing and H. Cruse, 'Stick insect locomotion in a complex environment: climbing over large gaps', *The Journal of Experimental Biology*, **207**, 1273–1286, (2004).

[5] J. Bongard, V. Zykov, and H. Lipson, 'Resilient machines through continuous self-modeling', *Science*, **314**(5802), 1118 – 1121, (November 2006).

[6] Rodney A. Brooks, 'Intelligence without reason', in *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI-91)*, eds., John Myopoulos and Ray Reiter, pp. 569–595, Sydney, Australia, (1991). Morgan Kaufmann publishers Inc.: San Mateo, CA, USA.

[7] H. Cruse, 'Feeling our body - the basis of cognition?', *Evolution and Cognition*, **5**(2), 162–173, (1999).

[8] A. Cully, J. Clune, D. Tarapore, and J.-B. Mouret, 'Robots that can adapt like animals', *Nature*, **521**(7553), 503–507, (May 2015).

[9] V. Gallese and G. Lakoff, 'The brain's concepts: the role of the sensory-motor system in conceptual knowledge', *Cognitive Neuropsychology*, **22**(3-4), 455–479, (June 2005).

[10] A. M. Glenberg, 'What memory is for', *Behavioral and Brain Sciences*, **20**(1), 1–55, (March 1997).

[11] G. Hesslow, 'Conscious thought as simulation of behaviour and perception', *Trends in Cognitive Sciences*, **6**(6), 242–247, (2002).

[12] D. McFarland and T. Bösser, *Intelligent behavior in animals and robots*, MIT Press, Cambridge, MA, 1993.

[13] J. Paskarbeit, J. Schmitz, M. Schilling, and A. Schneider, 'Layout and construction of a hexapod robot with increased mobility', in *IEEE/RAS-EMBS International Conference on Biomedical Robotics and Biomechatronics*, pp. 621–625, (2010).

[14] M. Schilling, 'Universally manipulable body models — dual quaternion representations in layered and dynamic MMCs', *Autonomous Robots*, **30**(4), 399–425, (2011).

[15] M. Schilling and H. Cruse, 'What's next: Recruitment of a grounded predictive body model for planning a robot's actions', *Frontiers in Psychology*, **3**(383), (2012).

[16] M. Schilling, H. Cruse, and P. Arena, 'Hexapod Walking: an expansion to Walknet dealing with leg amputations and force oscillations', *Biological Cybernetics*, **96**(3), 323–340, (2007).

[17] M. Schilling, T. Hoinville, J. Schmitz, and H Cruse, 'Walknet, a bio-inspired controller for hexapod walking', *Biological Cybernetics*, **107**(4), 397–419, (2013).

[18] M. Schilling, J. Paskarbeit, A. Hüffmeier, A. Schneider, J. Schmitz, and H. Cruse, 'A hexapod walker using a heterarchical architecture for action selection', *Frontiers in Computational Neuroscience*, **7**(126), (2013).

[19] M. Schilling, J. Paskarbeit, J. Schmitz, A. Schneider, and H Cruse, 'Grounding an internal body model of a hexapod walker – control of curve walking in a biological inspired robot—control of curve walking in a biological inspired robot', in *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2012*, pp. 2762–2768, (2012).

[20] Malte Schilling and Holk Cruse, 'Reacog, a minimal cognitive controller based on recruitment of reactive systems', *Frontiers in Neurorobotics*, **11**, 3, (2017).

[21] J. Schmitz, A. Schneider, M. Schilling, and H. Cruse, 'No need for a body model: Positive velocity feedback for the control of an 18-DOF robot walker', *Applied Bionics and Biomechanics, Special Issue on Biologically Inspired Robots*, **5**(3), 135–147, (2008).

[22] M. Schumm and H. Cruse, 'Control of swing movement: influences of differently shaped substrate', *Journal of Comparative Physiology [A]*, **192**(10), 1147–1164, (2006).

[23] L. Steels, 'Intelligence with representation', *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, **361**(1811), 2381–2395, (2003).

[24] M. Wilson, 'Six views of embodied cognition', *Psychonomic Bulletin & Review*, **9**(4), 625–636, (2002).