Using Sensorimotor Contingencies for Terrain Discrimination and Adaptive Walking Behavior in the Quadruped Robot Puppy

Matej Hoffmann[1], Nico M. Schmidt[1], Rolf Pfeifer[1], Andreas K. Engel[2], and Alexander Maye[2]

[1] Artificial Intelligence Laboratory, Department of Informatics, University of Zurich
Andreasstrasse 15, 8050 Zurich, Switzerland
{hoffmann,nschmidt,pfeifer}@ifi.uzh.ch
[2] University Medical Center Hamburg-Eppendorf
Dept. of Neurophysiology and Pathophysiology
Martinistr. 52, 20246 Hamburg, Germany
{a.maye,ak.engel}@uke.de

Abstract. In conventional "sense-think-act" control architectures, perception is reduced to a passive collection of sensory information, followed by a mapping onto a prestructured internal world model. For biological agents, Sensorimotor Contingency Theory (SMCT) posits that perception is not an isolated processing step, but is constituted by knowing and exercising the law-like relations between actions and resulting changes in sensory stimulation. We present a computational model of SMCT for controlling the behavior of a quadruped robot running on different terrains. Our experimental study demonstrates that: (i) Sensory-Motor Contingencies (SMC) provide better discrimination capabilities of environmental properties than conventional recognition from the sensory signals alone; (ii) discrimination is further improved by considering the action context on a longer time scale; (iii) the robot can utilize this knowledge to adapt its behavior for maximizing its stability.

Keywords: active perception, terrain recognition, object recognition, developmental robotics, adaptive behavior

1 Introduction

In the majority of approaches to robot control the extraction and classification of features from the sensory input is a crucial processing step that has a critical effect on the behavioral performance of the artificial agent. Ever more complex methods are employed to detect type and position of objects, to recognize landmarks and obstacles, or to infer the spatial configuration of the surrounding area. In mobile robotics, for example, this problem is typically solved by employing several distal (non-contact) sensors: cameras, laser range finders, and possibly also radar. Terrain classification into traversable vs. non-traversable is done in a supervised manner through a set of labeled terrain examples [1]. This is used to update an internal representation of the world – a 2D occupancy grid that in turn is used for planning a collision-free path. Although recent studies suggest that the traditional "sense-think-act" approaches can also be extended to real-world environments, their task domain is still limited.

The inherent problem of these approaches, in our view, is that they treat perception as a separate, *passive* process that is detached from the agent's actions. A "sensory snapshot" of the environment is taken that is then mapped onto the states of an internal world model. However, we believe that perception in biological agents has a different character. First, it is active. This view can be traced back to the pragmatic philosopher John Dewey [3], and it was later picked up by research in active perception (see [4] for an overview). Second, perception occurs through the body. The information that reaches the brain is thus critically shaped by the active generation of sensory stimuli and by the agent's embodiment (this is quantified in [10], for instance). Sensorimotor Contingency Theory (SMCT)[15, 14] as a representative of action-oriented approaches ascribes sensory awareness and perception to the exercise of knowledge about the lawful relations between actions and resulting changes in the sensory signals, called Sensory-Motor Contingencies (SMCs), instead of activating an internal representation of the perceived object.

We have recently developed a computational model for SMCs and demonstrated its application in an object-recognition task [11]. Here we apply the same model for controlling a robot with a completely different embodiment: a quadruped "dog" robot. We start by investigating how different gaits and terrains modulate the sensory information collected by the robot. Next we demonstrate that taking the action explicitly into account improves the terrain classification accuracy. Taking the context of longer sensorimotor sequences into account can further improve the classification performance. Finally, we show that the robot can successfully deploy its perception of the properties of different grounds to select gaits from a given repertoire to maximize its stability.

2 Related Work

The importance of sensorimotor information for object recognition in humans is evident from studies of neurological disorders [22], even though it is sometimes assigned only the role of a fall-back system [18]. In a scenario similar to ours, E.J. Gibson et al. [5] studied how infants perceive the traversability of the environment, implicitly taking into account their mode of locomotion – walking or crawling – and exploiting not only visual but also tactile information. In general, perceptual categorization in biological agents is a hard problem [7] resulting from a complex interplay of the brain, body and environment, and the individual effects are hard to separate. In this regard, robotics has provided efficient tools to test these effects independently.

First, Pfeifer and Scheier [16] have demonstrated how sensorimotor coordination can greatly simplify classification or categorization in a study where mobile robots distinguish between big and small cylinders by circling around them. Whereas this would be very difficult from a static camera picture when the distance to the object is not known, different angular velocities resulting from circling around them render the problem much easier. Similar results emerged from studies in artificial evolution: the fittest agents were those engaging in sensory-motor coordinated behavior [2].

Second, perception can be facilitated by the morphology of the body and the sensory apparatus (see examples in [8]). In legged robots that engage in different terrains, proprioceptive sensors can be particularly useful. In a previous study in our platform, we have shown how information regarding the robot's position and orientation can be extracted [17]. A combination of proprioceptive sensors has been successfully employed in a terrain recognition task in a hexapod [6].

Third, the action that caused a sensory stimulation can be explicitly taken into account in a classification task. This has been done in [19], where sensory data resulting from different actions are clustered separately. In [20], traversability categories are predefined and the robot learns – for each action separately – a mapping from initial percepts to these categories.

Many more approaches employ some form of sensorimotor information, but to our knowledge the approach we will present here is one of the few in that actions play a constitutive role for the perception of the agent as proposed by SMCT. Our method allows for a context given by the sequence of previous actions, and it is inherently multimodal. In addition, we will test the hypothesis that longer sensorimotor sequences are needed for object categorization (i.e., the ground the robot is running on in our case). Furthermore, to demonstrate the behavioral relevance of the classification capabilities for the agent, we present a closed-loop system that employs the perception of the properties of different grounds to select gaits from a given repertoire to maximize stability.

3 Methods and experiments

3.1 Robot and Experimental Setup

The Puppy robot (see Fig. 1 left) has four identical legs driven by position-controlled servomotors in the hips. It has passive compliant joints at the knees. We prepared five sets of position control commands for the servomotors, resulting in five distinct gaits (bound forwards, bound left/right, crawl, trot backwards), each of them with a periodic motor signal at 1 Hz. Four potentiometers measured the joint angles on the passive knee joints, and 4 pressure sensors recorded forces applied to the robot's feet. Linear accelerations (in X, Y, and Z direction) were measured by an onboard accelerometer. In total we used 11 sensory channels, jointly sampled at 50Hz.

To investigate the long-term properties of our approach, we additionally designed a model of Puppy in Webots [21], a physics-based simulator (see Fig. 1 right). For this model we used the same gait repertoire (2 gaits had to be adapted) plus 4 additional gaits (turn left/right, pace, walk), obtaining a repertoire of nine gaits. In both cases, gaits (actions) were exercised in 2-second-intervals during which the sensory data were collected, forming sensorimotor epochs of 2 seconds. At the end of each epoch the robot could change the gait.

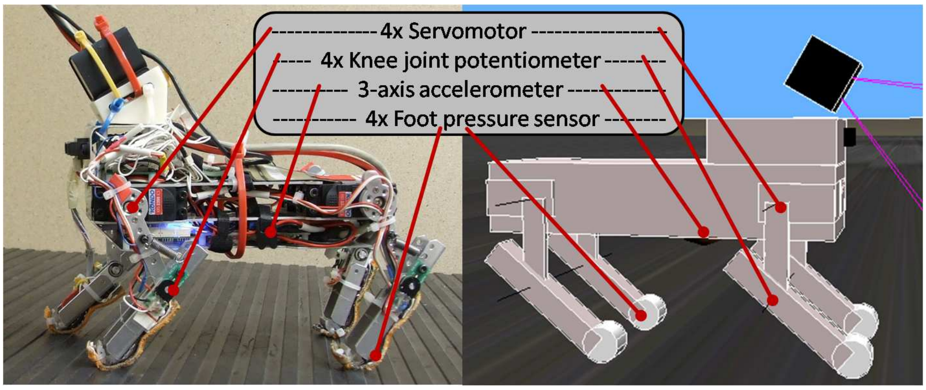

Fig. 1. Real and simulated robot and the sensor suite. The robot is 20 cm long. The camera and infrared sensors that are also mounted on the robot were not used in the experiments.

For the real robot, we prepared a small wall-enclosed arena of 2x1 m. Four different ground substrates covered the ground: plastic foil, cardboard, Styrofoam and rubber. These materials differed in friction and also in structure (cardboard and rubber had ridges). In the simulator, the arena was much bigger in size (25x25 m), so encounters with the walls were much less frequent. The "foil", "cardboard", and "rubber" were flat but differed in Coulomb friction coefficients (μ = 2, 11, and 20 respectively). To increase the differences between the substrates in the simulator, the "Styrofoam" ground ($\mu = 9$) was made uneven with randomly placed smooth bumps of up to 3 cm height.

3.2 Feature Computation

For effective processing of sensorimotor information we compressed the raw data by extracting some simple features. For the action space we chose a high abstraction level and used the gait as a single feature. In the sensory space, following a similar strategy as we used in [17], we took advantage of the periodic nature of the locomotion and created period-based features as follows: (1) sum of knee amplitudes of all legs in a period,[3] (2) sum of standard deviations of all knee joints, (3) sum of mean pressures in each foot, (4) sum of standard deviations of each foot pressure signal, (5-7) mean accelerations along X,Y, and Z-axis respectively, (8-10) standard deviations of the accelerometer signals. Since frequent

[3] Note that the knees are passive compliant.

gait transitions disrupt the locomotion and impact also the sensory values, only the last second (i.e. the second locomotion period) from each 2s epoch was used for the feature computation. Continuous feature values were used for classification (Section 4.1); for learning the sensorimotor contingencies and optimizing the behavior using a Markov model (Section 4.2), each feature was quantized to two levels only.

3.3 A Markov Model of SMCs

We employed the model that we presented in [11, 12] with the necessary adaptations to the Puppy robot. The basic idea is to consider actions and resulting changes in sensory signals in an integrated manner, and to keep a record of sequences of actions and sensory observations. For each epoch, the action a (the gait in this case) and a vector of n sensory features observed during execution of a are concatenated to a single vector $ao(t) = [as_1 s_2 \dots s_n]$ that we call an action-observation pair. Based on the sequence of action-observation tuples that the robot experiences over time $c^h = [ao(t), ao(t-1), \dots ao(t-h)]$, the model samples the conditional probability distributions $P^h(ao(t+1)|c^h(t))$, i.e. the probability of experiencing a particular action-observation pair in the next time step given a finite history h of previous pairs. In this study we use $h = 0 \dots 4$. This probability distribution is what we call the extended Sensori-Motor Contingencies (eSMC) of an agent, and a particular combination of $ao(t+1)$ and $c^h(t)$ is a specific sample that in addition to its probability of occurrence can have other properties like a value.

3.4 Value System and Action Selection

We extended the basic idea of SMCT by a value system and an action selection algorithm. For each epoch t, we define the value[4] of the robot's state by a weighted sum of three components:

$$v(t) = -tumbled - 0.4regularity - 0.1speed$$

We used the signal of the accelerometer in Z direction to determine if the robot is upright ($tumbled = 0$) or has tipped over ($tumbled = 1$). The similarity of the sensory patterns at the knee joints between the first and second period during an epoch is reflected in the *regularity* value (1 for identical patterns during both periods), and the normalized velocity computed from the robot's global coordinates yields the *speed* value.

We have devised a stochastic action selection algorithm that attempts to optimize the temporal average of the internal value. It selects actions that have shown to activate eSMCs with high internal values, and explores the consequences of new actions when no or only bad prior experiences exist in a given

[4] In reinforcement learning terms, this would be called reward - it is the immediate reward signal associated with each state.

situation. For each action-observation sequence $c^h(t)$ a record of actions executed next $a_{next}(c^h(t))$ and the average value $v(a_{next}(c^h(t)) = \sum_n v(t+1)/n$ is kept, where n is the number that action a_{next} was executed when context $c^h(t)$ was encountered, and $v(t+1)$ is the resulting value. Different history lengths h may yield different value information. Since we consider longer matches between a particular action-observation sequence and the stored eSMCs as a more accurate estimation of the state, preference is given to the value information from longer matching histories. When the robot later experiences the same context again, it knows the average value of the actions it has tried before. Random values get assigned to the other actions. To avoid a predominantly random exploration in the initial learning phase when the robot has only little sensorimotor knowledge, the expected value for the most recently executed action is given by the internal value of the last epoch. This favors the continuation of successful actions, and switching to another action otherwise. The action with the highest expected value $\hat{a} = \arg\max_a v(a_{next}(c^h(t))$ is then executed with a probability $p(\hat{a}) = v(\hat{a}) + 1$.

4 Results

4.1 Perception and Discrimination of Different Grounds

In this section, we want to quantitatively assess the effect of considering actions and the resulting changes in sensory stimulation in an integrated manner. First, we compare the respective influence of the action (the gait the robot is running with) and the environment on the sensory data. Second, focusing on the ground discrimination, we demonstrate how explicitly incorporating the action that has induced a sensory stimulation improves the environment classification. Finally, we study the effect of longer sensorimotor sequences, testing our hypothesis that these are required for object categorization, whereby, from the robot's perspective, different grounds correspond to different objects in our scenario.

We have collected data from the real (4 x 20 minutes, i.e. 4 x 600 epochs) and simulated version of the robot (4 x 4 hours, i.e. 4 x 7200 epochs) running separately on the different substrates. After every epoch a new action was chosen at random. If the robot tumbled, it was manually (real robot) or automatically (simulator) returned to an upright position at the same location and two epochs following this event were discarded. A reflex for backing up from the walls was built in. Epoches when the robot was backing up (frequent in the real robot) were not discarded but entered the regular learning process. A naïve Bayes classifier (diagonal covariance matrix estimate, stratified 10-fold cross-validation) was trained to classify either the action or the ground substrate given the sensory observations and actions during the previous epochs.

Ground and Gait Discrimination from Sensory Data Only. To assess the dependencies of the sensory signals from the gait or ground, respectively, we collapsed the data across gaits (for assessing ground effects) or across grounds (for assessing gait effects). In the real Puppy, the classifier determined the correct gait from the sensory data in 72.4% of the cases, and in 81.6% in the simulation. In contrast, the ground recognition rates were lower, 67.2% for the real Puppy and 43.1% in the simulation (see also Fig. 2, top-most bars). This shows that gaits and grounds have a similarly strong effect on the sensory patterns in the real robot. In the simulation the different materials induce similar sensory patterns and hence, are difficult to distinguish. These figures serve as a baseline when we consider the classification of joint action and sensor information next.

Ground Discrimination Using Action Information. We separated the data into sets for each gait and classified the grounds on each set individually. Afterwards we averaged the ground recognition rate over all gaits. In comparison to the ground recognition using a single classifier, the action-dependent classification schema reaches an improved accuracy of 75.7% for the real robot. Considering only the gait yielding the best recognition rate, this value increases to 80.2%. In the simulation this increase is even more pronounced, from 43.1% to 62.9% and 78.3%, respectively (see Fig. 2, second bars from top). This indicates that taking the action that caused a sensory observation into account is more specific for the environmental condition than analyzing the sensory data alone.

Ground Discrimination Using Action Sequences. The sensorimotor patterns induced by a single action may often be similar even if the agent interacts with different objects. As suggested by SMCT, longer sequences of interaction with an object may be needed in order for the object to leave a unique "footprint". We confirmed this hypothesis by splitting the data further into sets for specific sequences of 2 or 3 consecutive actions, and averaging again over all sequences. The sensory feature vectors from consecutive epochs were concatenated. For a sequence of two gaits, the ground classification accuracy rises to 84.7% in the real robot, and to 70.6% in the simulation. Considering a sequence of 3 gaits further improves accuracy (see Fig. 2). Here, the gait sequence-specific classifier with the highest accuracy achieves a 100% recognition rate. This means that the sensorimotor patterns of this action sequence are apt for a reliable recognition of the different grounds.

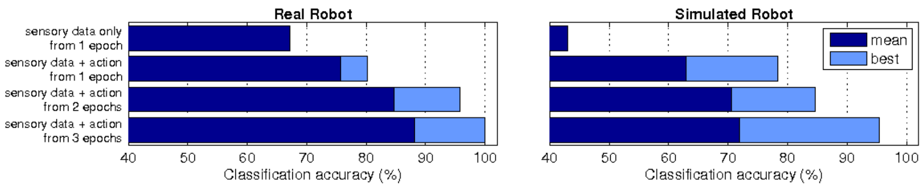

Fig. 2. Comparison of the ground classification accuracies when the action context is taken into account to different degrees. (left) Real robot. (right) Simulated robot.

4.2 Selecting Gaits to Optimize Behavior

Next we want to demonstrate how the better discrimination capabilities that result when a longer action context is considered can be used by the robot to improve its behavior. We let the simulated robot run on 4 different grounds, and used the Markov model (Sec. 3.3) to learn eSMCs for the 9 gaits from its repertoire. Each eSMC had an associated value given by the value function described in section 3.4.

With progressing sensorimotor knowledge, the robot preferred to choose gaits that improved its internal value, providing swift, smooth and stable locomotion. The plots of the value function in Fig. 3 show that a basic set of gaits that "feel good" to Puppy (i.e. maximize the value function) is found after only about 1.000 epochs (around 8 minutes). On cardboard it takes more than 2.000 epochs to arrive at a reasonable gait combination. Afterwards the robot tries to further improve its behavior by selecting from these comfortable gaits with different probabilities. As one would expect, the optimal gait sequence depends on the material properties of the grounds. Except for the plastic foil, Puppy prefers a mixture of walking back and turning left or right. It is most successful in epochs when it reduces the frequency of turns in favor of walking back. On plastic foil, the most successful gait is pacing, while turning left seems to be a less favorable gait. On cardboard, turning left is selected more frequently than turning right, though, while on rubber both turning actions are chosen with about the same frequency.

On the rough styrofoam, the value function is dominated by frequent tipping of the robot. Compared to the three flat grounds the value remains at a low level, and the separation into favorable and unpleasant gaits is less pronounced. The order of preference seems to be maintained, though.

The improvement of the internal value is not monotonic, but proceeds in a rather oscillatory manner. Intervals in which the robot had sufficient sensorimotor knowledge to optimize its behavior alternated with epochs in which it learned new eSMCs. With the sensorimotor knowledge growing, episodes with optimal behavior become more frequent and last longer. On cardboard, for example, behaviors that maximize the value function are found after about $2 \cdot 10^4$ epochs, and the increasing width of the peaks in the value function indicate that the robot spends more and more time in these optimal behaviors. A similar observation can be made on plastic foil. On rubber, the knowledge about favorable behavior around $2 \cdot 10^4$ seems to be lost afterwards, but it can be expected that the exploration process leads to a further improvement beyond the analyzed interval. Since the value function was designed to never reach zero, corresponding to a state of perfect harmony, the robot keeps on exploring the potential to further improve its fitness.

5 Conclusion and Future Work

In this study we have investigated sensorimotor classification of different substrates in a quadruped robot from the perspective of SMCT. First, we have

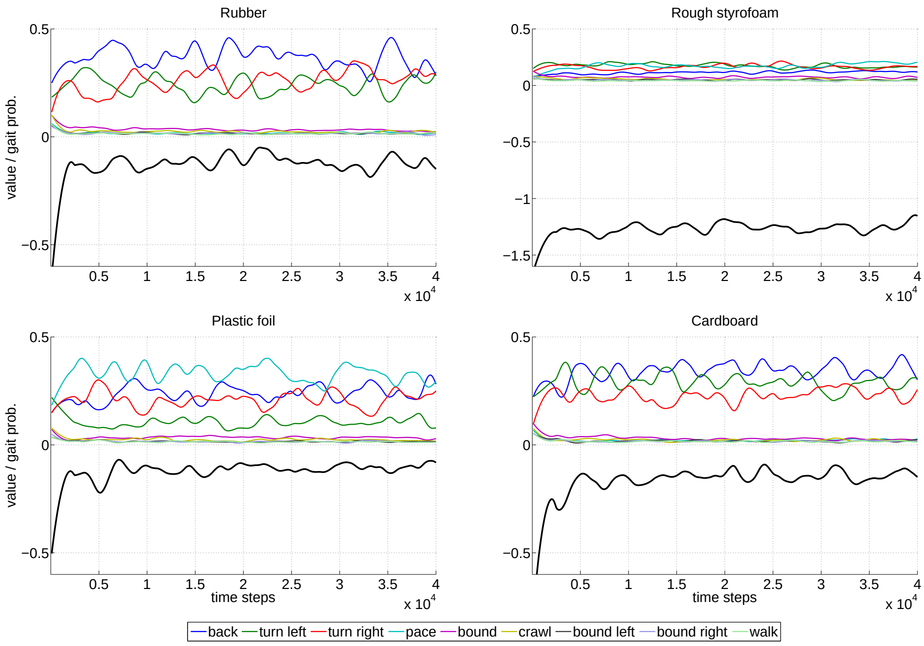

Fig. 3. Value (black curve under the abscissa) and gait selection frequencies (above) over time on 4 ground substrates (data from simulator). All curves have been smoothed with a weighted linear least squares and a 2nd degree polynomial model in a moving window of 5.000 samples. To appreciate the time course of the value function, the initially low values have been clipped. Note the different scale of the value function for rough styrofoam.

demonstrated how sensory stimulation patterns critically depend on the actions the robot is exercising. If the robot wants to recognize the object or environment it is interacting with, like the terrain type in our case, the action (gait) that gives rise to the experienced sensory stimulation needs to be considered. In addition we have shown that deployment of longer action contexts further improves the discrimination capabilities. Our approach demonstrates that the robot successfully engages the acquired sensorimotor knowledge to optimize its behavior by selecting appropriate gaits on different ground substrates.

Apart from serving as a model of SMCT, our work has also substantial application potential. Autonomous, perception-based, off-road navigation is a hot research topic in mobile robotics (e.g., [9]). Unlike traditional approaches that rely on passive long-distance perception using high resolution sensors, we have hinted at the potential of a radically different approach: terrain perception through active generation of sensory stimulation in a multimodal collection of low-resolution sensors (for learning eSMCs, 1 bit per sensory channel was used). Taking action-observation sequences into account and exploiting the robot's rich body dynamics to simplify the structure of the sensory information, an advantageous transformation of the input space for classification can be achieved.

In the current study we have employed only proprioceptive and contact sensors. These have proven very effective in ground discrimination and, in conjunction with a simple one-step prediction of the best next action based on the current sensorimotor context, the robot could optimize its behavior. However, these sensors provide little information about the terrain beyond the robot's current location. Distal sensors (like infrared or vision), on the other hand, could provide information about future events that could likewise be exploited for perceptual categorization and further improvement of the behavior. A promising approach in this respect uses internal simulation in sensorimotor space to find action sequences that optimize the success of the agent with a longer temporal horizon [13, 19]. Alternatively, reinforcement learning algorithms could be employed. Traversability in general may be a suitable touchstone to compare different approaches to use sensorimotor information for controlling robots. This will be the direction of our future work.

Acknowledgement

This work was supported by the Sino-German research training group CINACS, DFG GRK 1247, `www.cinacs.org`, and by the EU project 'Extending sensorimotor contingencies to cognition - eSMCs', FP7-ICT-270212, `esmcs.eu`.

References

1. J.A. Bagnell, D. Bradley, D. Silver, B. Sofman, and A. Stentz. Learning for autonomous navigation. *Robotics & Automation Magazine, IEEE*, 17 (2):74–84, 2010.
2. R.D. Beer. The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11:209–243, 2003.
3. J. Dewey. The reflex arc concept in psychology. *Psychological Review*, 3:357–370, 1896.
4. A.K. Engel. Directive minds: how dynamics shapes cognition. In Stewart J., Gapenne O., and Di Paolo E. A., editors, *Enaction: Towards a New Paradigm for Cognitive Science*, pages 219–243. MIT Press, Cambridge, MA, 2011.
5. E.J. Gibson, G. Riccio, M.A. Schmuckler, T.A. Stoffregen, D. Rosenberg, and J. Taromina. Detection of traversability of surfaces by crawling and walking infants. *Journal of Experimental Psychology*, 13(4):533–544, 1987.
6. P. Giguere and G. Dudek. Clustering sensor data for autonomous terrain identification using time-dependency. *Autonomous Robots*, 26:171–186, 2009.
7. S. Harnad. *Handbook of categorization in cognitive science*, chapter Cognition is categorization. Elsevier, 2005.
8. M. Hoffmann and R. Pfeifer. *The Implications of Embodiment: Cognition and Communication*, chapter The implications of embodiment for behavior and cognition: animal and robotic case studies, pages 31–58. Imprint Academic, 2011.
9. L.D. Jackel, E. Krotkov, M. Perschbacher, J. Pippine, and C. Sullivan. The DARPA LAGR program: goals, challenges, methodology, and phase I results. *Journal of Field Robotics*, 23(11/12):945–973, 2006.

10. M. Lungarella and O. Sporns. Mapping information flow in sensorimotor networks. *PLoS Computational Biology*, 2 e144(10):1301–1312, 2006.
11. A. Maye and A.K. Engel. A discrete computational model of sensorimotor contingencies for object perception and control of behavior. In *2011 IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 3810–3815. IEEE, May 2011.
12. A. Maye and A.K. Engel. Time scales of sensorimotor contingencies. In H. et al. Zhang, editor, *Brain-Inspired Cognitive Systems 2012*, volume 7366 of *LNAI*, pages 240–249. Springer-Verlag Berlin, Heidelberg, 2012.
13. R. Möller and W. Schenck. Bootstrapping cognition from behavior – a computerized thought experiment. *Cognitive Science*, 32(3):504–542, 2008.
14. A. Noë. *Action in perception.* MIT Press, 2004.
15. J.K. O'Regan and A. Noë. A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24:939–1031, 2001.
16. R. Pfeifer and C. Scheier. Sensory-motor coordination: The metaphor and beyond. *Robotics and Autonomous Systems*, 20:157–178, 1997.
17. M. Reinstein and M. Hoffmann. Dead reckoning in a dynamic quadruped robot: Inertial navigation system aided by a legged odometer. In *Robotics and Automation (ICRA), 2011 IEEE Int. Conf. on*, pages 617–624, 2011.
18. A. Sirigu, J. R. Duhamel, and M. Poncet. The role of sensorimotor experience in object recognition. A case of multimodal agnosia. *Brain : a journal of neurology*, 114 (Pt 6):2555–2573, December 1991.
19. E. Ugur, E. Oztop, and E. Sahin. Goal emulation and planning in perceptual space using learned affordances. *Robotics and Autonomous Systems*, 59:580–595, 2011.
20. E. Ugur and E. Sahin. Traversability: a case study for learning and perceiving affordances in robots. *Adaptive Behavior*, 18:258–284, 2010.
21. Webots. www.cyberbotics.com. Commercial Mobile Robot Simulation Software.
22. D.A. Wolk, H.B. Coslett, and G. Glosser. The role of sensory-motor information in object recognition: Evidence from category-specific visual agnosia. *Brain and Language*, 94(2):131 – 146, 2005.