Using Kinematically Complex Robots for Case Studies in Embodied Cognition

Y. Kassahun[†], M. Edgington[†], J. de Gea[†], E. Kirchner[†], D. Spenneberg[*] and F. Kirchner[*†]

[†]University of Bremen, Faculty of Mathematics and Computer Science,
Robert Hooke Str. 5, 28359 Bremen, Germany
[*]DFKI - German Research Center for Artificial Intelligence, Robotics Lab,
Robert Hooke Str. 5, 28359 Bremen, Germany
kassahun@informatik.uni-bremen.de

Abstract— **We present two case studies in embodied cognition which use kinematically complex robots for spatial cognition and concept forming. The first case study involves substrate classification on the basis of primarily proprioceptive data. During walking over various substrates a legged robot generates certain substrate specific sensory motor patterns. The acquired data is used for training a growing self-organizing neural network, which is connected with a standard output layer representing different substrates. The second case study is concerned with a recognition system which learns to recognize objects based on multimodal sensorimotor coordination. The sensorimotor coordination is generated through interaction with the environment. The system uses a learning architecture which is composed of reactive and deliberative layers. The reactive layer consists of a database of behaviors that are modulated to produce a desired behavior. We have implemented in the learning architecture an object manipulation behavior inspired by the concept that infants learn about their environment through manipulation [1]. While manipulating objects, the agent records both proprioceptive data and exteroceptive data. Both of these types of data are combined and statistically analyzed in order to extract important parameters that distinctively describe the object being manipulated. This data is then clustered using the standard k-means algorithm and the resulting clusters are labeled. The labeling is used to train a radial basis function network for classifying the clusters. It has been found that the trained neural network is able to classify objects even when only partial sensory data is available to the system. Our preliminary results in both case studies demonstrate that kinematically complex robots are suitable for learning about their environment from experience and provide a new useful class of proprioceptive information in contrast to wheeled systems.**

Keywords— **Substrate classification, Learning through interaction, Sensorimotor coordination, Object recognition**

I. Introduction

This article describes work in the field of "Embodied AI"[2]. In Embodied AI the features of the body of an agent are granted a major role in the process of forming concepts of the experienced world. Only through the abilities of its body can an agent be 'situated' and 'embodied' in the world. The representations it uses to represent the world must be based on its own abilities and experiences instead on models given by a human developer. In recent years, the concept of sensorimotor coordination has gained increasing attention in the psychological and artificial intelligence communities, and it is considered by many to be a prerequisite for developing higher levels of cognition in intelligent beings [3], [4]. For example, classifying an object often requires that one manipulates objects within one's environment. In this work, motor skills are coordinated with sensory information (tactile, visual, etc.) in order to better identify an object or the environment. It has been shown that sensorimotor coordination can be exploited in solving classification problems [5], [6], [7], [8]. One should note, however, that this process exploits not only the motor actions of a body, but also the intrinsic structure of the body itself. The body structure of a being or a machine enforces certain constraints on the allowable range of motions by providing natural configurations, in which energy consumption and stress are minimized. Therefore, the body plays an important role in structuring the sensorimotor coordination. One shortcoming of most existing methods that use sensorimotor coordination is that they are able to recognize only a limited number of objects. Additionally, most existing methods are difficult to extend. The typical application of such methods is to recognize idealized objects, and the testing of the methods is usually done only in simulation or on simple robotic platforms. Furthermore, their ability to scale when used on complex robots has not been experimentally verified. From our viewpoint, the reasons for the shortcomings of the existing methods are twofold:

1. At present, there is no firm theoretical framework for studying correlations within and between sensorimotor modalities for object recognition tasks. Very few approaches apply statistical and information theoretic analyses to study the sensorimotor coordination of data taken from real robots [9].

2. Kinematically complex robots capable of increasing the role of the body in the process of learning and recognition are not commonly used. Most of the time wheeled robots with few degrees of freedom or simulated robotic arms are used as test beds.

To pursue the "Embodied AI" approach in the field of robotics, many researchers suggest the development of systems with richer system-environment interaction. Therefore, in the last years we focused our research on complex legged robots which possess a rich repertoire of sensor and motor abilities, especially proprioceptive sensing. In this article we present two case studies, which investigate the usefulness of kinematically complex robots in studying spatial recognition and in concept forming.

II. Case Study I: Using proprioceptive data for classification of terrain structures

Initial experiments to use proprioceptive data for spatial categorization have been carried out using the 8-

legged robot SCORPION [10] (see Fig. 1), which provides a multitude of proprioceptive sensor signals (motor current, joint position, pressure applied to the feet, pitch and roll). With this robot we conducted ex-

Fig. 1. The SCORPION Robot in the Experimental Test Bed

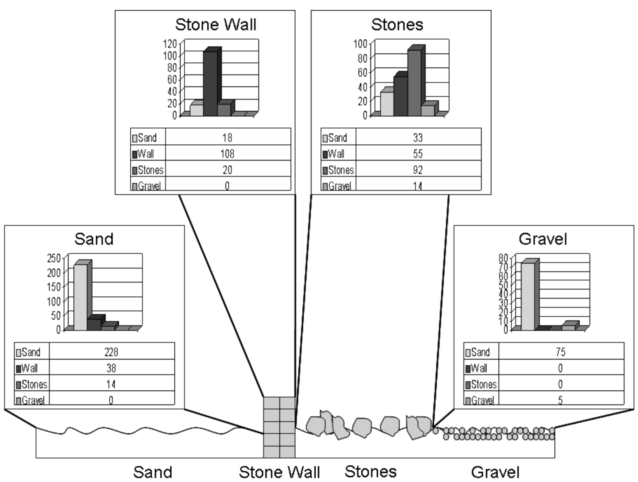

Fig. 2. Classification Results in the Test Bed

periments to investigate the possibility of using only proprioceptive data to generate information about the environment. By interacting with the environment, e.g., walking over it, the SCORPION produces a multitude of proprioceptive data. In this case study, we have investigated the feasibility of using this data to detect on which substrate the robot is walking. This ability can be especially useful for self-localization and mapping. The experimental platform is an indoor test bed, which contains different substrates: sand, a small stone wall, a rockfield and a gravel field. During several 2-minute-runs through the whole test bed (starting in the sand and ending in the gravel) all proprioceptive data was collected. In addition markers are manulally placed in the logged data every time the robot crosses into the next substrate, e.g., when the robot is at the border between the sand area and the stone wall. These resulting sets of data pairs are then used as base material for a training set for a GCS classifier [11]. Growing cell structures are a vector based self organizing neural network approach. We were able to achieve promising results (see Fig. 2), which prove that proprioceptive data produced through interaction with the environment can be used for identifying substrates. The system was able to distinguish beteen three classes of subtrates. The rock field and the stone wall were easily distinguished from the other subtrates, however, sand and gravel were grouped into a single class due to the fact that for the SCORPION robot sand and gravel have quite the same properties. Thus from the perspective of the SCORPION robot these two different categories do not exist; in its world they would be treated as the same. To investigate this further we plan to develop an unsupervised learning method, which would allow the generation of categories that represent the robots view on the world. More details on these experiments can be found in [12].

III. Case Study II: Multimodal Object recognition based on proprioceptive and exteroceptive data

In the second case study [1], we present an extensible embodied object recognition system that can be used in complex real robots that learn through interaction with the environment. The system can be easily extended to use new object-features which distinctively describe the relevant characteristics of an object to be recognized. In this case study, we first give a short overview of the learning architecture we have used in order to implement object recognition through manipulation. We then explain the manipulation behavior and the recognition method used. Finally, we describe our experimental scenario and the results obtained.

A. Learning Architecture

The architecture we have adopted, shown in Figure 3, is a hybrid architecture which integrates a reactive system with a higher-level deliberative system. It is a biologically inspired learning system that is suitable for controlling and integrating spatial learning and representation techniques in mobile robots, allowing them to explore and navigate in unknown environments.

The sensory inputs first enter the **sensory perception processor** and textbfsituation analysis modules. The output of the **sensory perception processor** goes to the **reflex transfer**, **attentional transfer**, **habit transfer** and **goal directed transfer** modules. The **attentional transfer** module receives additional inputs from the **situation analysis** module. The output of the **reflex transfer** module goes to the **reflex motor program library** module. After being modulated by the **goal directed transfer** module, the outputs of **attentional transfer** and **habit transfer** go to the **motor program library**. Finally the signal from the **motor program library** or **reflex motor**

[1]sponsored by the DFG (SFB-TR8)

program library goes to the **low level motor control** unit.

A.1 Reactive System

The reactive system includes the **sensory perception processor**, **situation analysis**, **reflex transfer**, **attentional transfer**, **habit transfer**, **reflex motor program library** and **motor program library** modules of Figure 3.

Sensory Perception Processor

Proprioceptive and exteroceptive data produced through interaction with the environment are processed in the **sensory perception processor** module. This module consists of four subsystems (see Figure 4), which work together in order to learn, distinguish and identify different hypotheses that are relevant to the system. The figure shows an implementation of this module that is used in the experiment presented in this case study.

For the experiment, an unsupervised classifier system is implemented that considers both proprioceptive and exteroceptive data, represented by the vectors $\vec{S}_P$ and $\vec{S}_E$, respectively. This classifier identifies and labels clusters of similar sensorimotor stimuli together. Additionally, it generates a cluster probability $\vec{P}_{C,sm}(\vec{S}_P, \vec{S}_E)$. Each element of this vector represents a probability estimate of the likelihood that a set of sensorimotor inputs belongs to a labeled cluster. The cluster probability $\vec{P}_{C,sm}$ is then used to train the two remaining classifiers which classify based only on the exteroceptive sensory data vector $\vec{S}_E$ or on the proprioceptive data vector $\vec{S}_P$. These classifiers also generate $\vec{P}_C$ estimates, $\vec{P}_{C,ext}$ and $\vec{P}_{C,prop}$, which are combined with $\vec{P}_{C,sm}$ to generate an overall cluster probability $\vec{P}_{C,input}$.

Situation Analysis

This module is responsible for analyzing the proprioceptive and exteroceptive data in order to determine if some particular event has happened. This information is used by the **attentional transfer** module. Examples of events detected by this module are movements, changes in the sensory data and so on.

Reflex Transfer Module

This system tries to prevent the robot from being damaged. Depending on the output of the **sensory perception processor**, or on the raw sensory data, it activates certain reflex motor programs to protect the robot from danger. The module inhibits all the activities of the **habit transfer** and **attentional transfer** modules whenever it is active.

Attentional Transfer Module

This module orients the system to a given location based on the output of the **situation analysis** and **sensory perception processor** modules. An example could be to orient the system to a location where something moves, or where something interesting for the system is recognized.

Habit Transfer Module

This module is a collection of learned behaviors which have proven to be important to the system. The learned behaviors are generated in the **goal directed transfer** module and installed in the **habit transfer** module. If a behavior in the **habit transfer** module is not used for a long time, it will be deactivated and its importance will decrease over time.

B. The Goal Directed System

This system is responsible for high-level processing of hypotheses probabilities, and it is in this system that the world model and body model are generated through learning. Positive and negative rewards combined with the output of the **sensory perception processor** module are used to learn the world model and body model of the robot itself. The **goal directed transfer** module is also responsible for optimizing existing habits or adding a new habit. In addition to this, this module pre-modulates the output of the **habit transfer** and **attentional transfer** modules to affect which motor-programs will be active, and post-modulates the output of the `motor program library`, modifying the properties of the motor-programs that are currently active.

Biological Support

There is no doubt that a reward system is needed to be able to learn or reinforce behavior. Reinforcement can be described as a process where certain events (i.e., reinforcing stimuli) increase the probability of the behaviors they are associated with (e.g. [13]) while reward refers to the tendency of certain events to direct behavior, specifically, to elicit and reinforce approach behavior. In biological systems rewarding stimulation directly activates dopamine neurons (see [14]). Two major dopamine systems have been implicated in a wide variety of behavioral actions including locomotor activity. Those systems are the nigrostriatal and the mesolimbic/mesocortical dopamine systems. Numerous attempts have been made to develop an integrative theoretical model (e.g. [15], [16], [17], [18]). Studies which show the importance of a reward system revealed that electrical brain stimulation can control behavior in much the same fashion as conventional rewards (see [19]). They have also shown that electrical stimulation can activate reward systems involved in the control of natural behavior. Details can be found in [20].

In the architecture we have adopted, reward is the essential part of of the goal directed behavior. When the **goal directed system** gets certain sensory input from the **sensory perception processor**, it will produce certain behavior depending on the expected reward it receives. The reward and especially the *lack* of reward gives feedback to the system at to whether the behavior was or was not appropriate given the situation 'described' by the **sensory perception proces-**

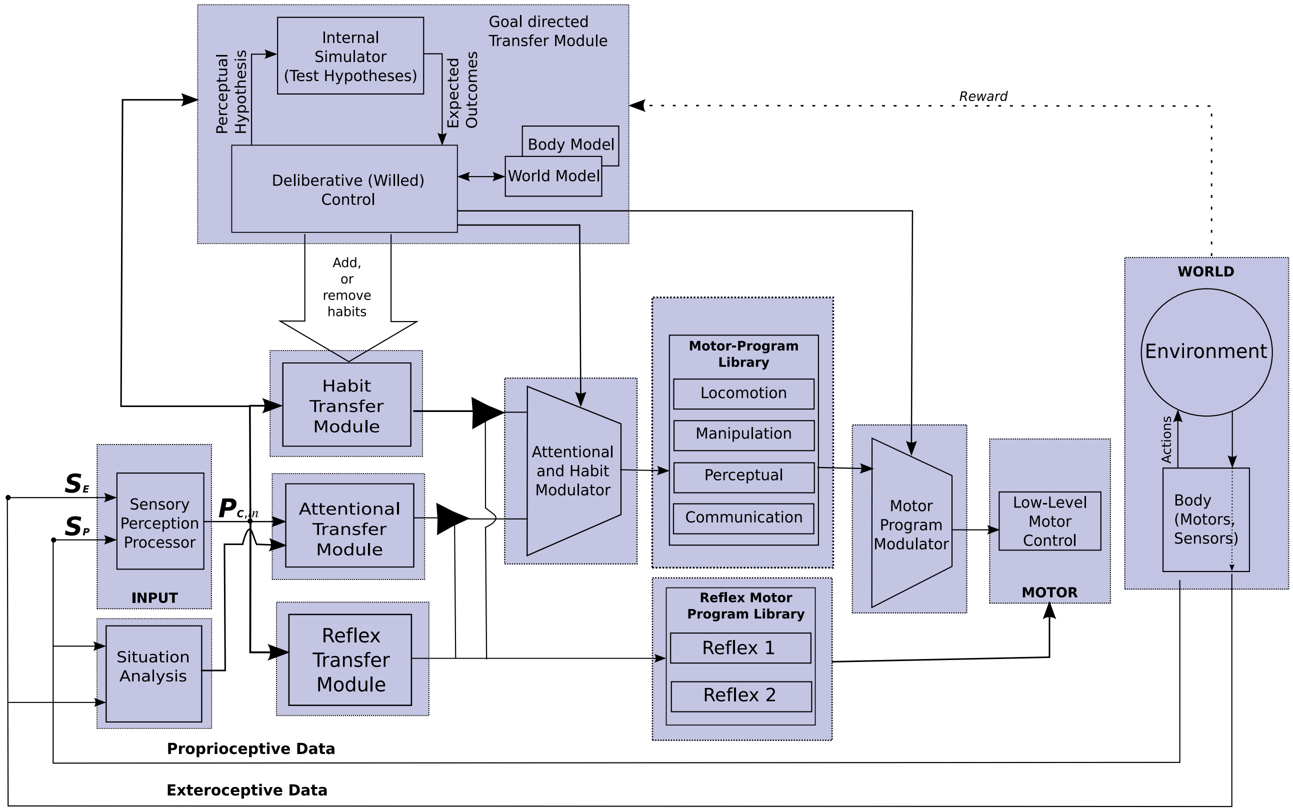

Fig. 3. System architecture. The shaded triangles in the figure show switches, which are controlled by the **reflex transfer** module. Whenever the **reflex transfer** module is active, the outputs of **habit transfer** and **attentional transfer** modules are inhibited.

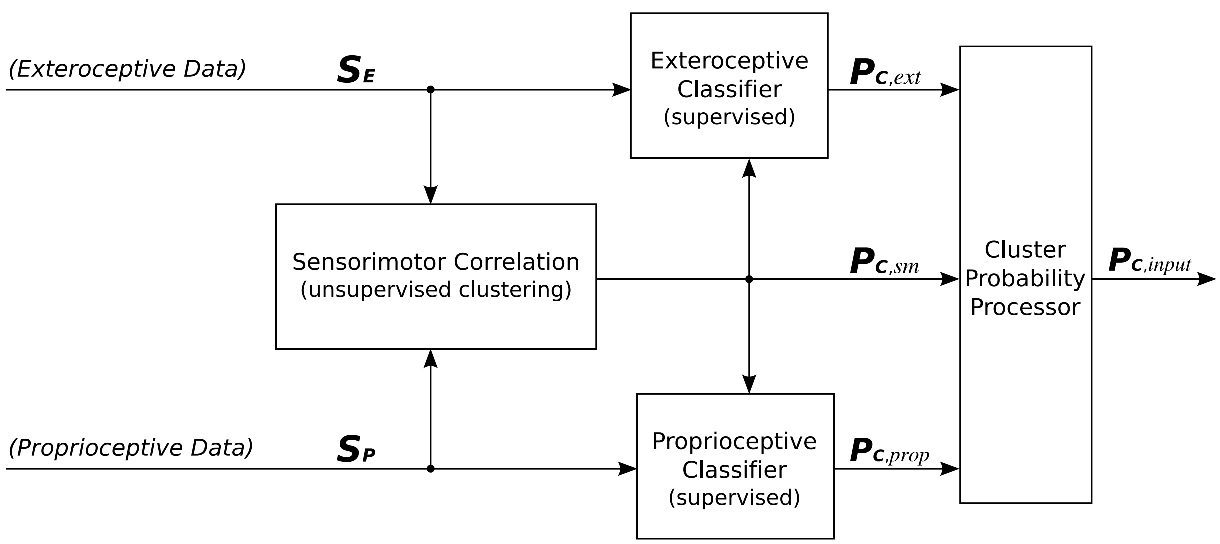

Fig. 4. The `sensory perception processor` module.

sor. Therefore, by learning about the environment, the system builds a model about the world. This model in turn can be used to simulate the likelihood of a reward following a certain behavior, given input from the **sensory perception processor**. A reward will consolidate the part of the model used and due to that reinforce the dominant behavior under repeated or very similar input. In case of lack of reward after the dominant behavior, the model that was used to predict the "best" behavior will be changed according to the outcome. In goal directed behavior, reward and *lack* of reward lead to the creation of a world model that connects certain input to the **sensory perception processor** with likely rewarded behavior. The manifestation of such a behavior is made possible by the activation and recombination of certain motor programs in the

motor program library and by creating new motor programs. Thus, certain motor programs are activated depending on the presence of specific inputs from the **sensory perception processor**.

Before a habit is added to the **habit transfer** module, it is learned in the goal directed system based on the reward received by the system. The **goal directed transfer** module uses a model based reinforcement learning similar to the Dyna system [21], where the model of the world is learned along with a habit. The model of the world predicts the next state and reward based on the current state and believed action. If the expected reward is missing under a given condition, the system will be "surprised" and will deactivate the current habit and shift to a goal directed learning procedure. During this time, a new habit may emerge and be installed in the **habit transfer** module.

C. The Recognition System

For the recognition system, we have not fully exploited the proposed learning architecture. The system described in this section mainly uses the **sensory perception processor** and **motor program library** modules. The embodied recognition system functions by manipulating objects in order to determine their specific characteristics. A manipulation motor-program has been implemented, added to the **motor program library**module, and made active for the experiment. This motor-program uses a potential field method [22] to generate a trajectory for an end-effector to reach an object. The basic idea is to create a mathematical description of a virtual potential field acting within the workspace of the manipulator. Regions in the workspace that are to be avoided are modelled by repulsive potentials (energy peaks) and the target region/point is modelled by an attractive potential (energy valley). The sum of repulsive and attractive potentials provides a representation of the workspace topology. By following the gradient (i.e. the minimum potential field at each step), a path towards the goal is generated. One fundamental difference between this method and classical path planning is that here "planning" is not done in the usual sense. Rather, a path is incrementally computed that will end at the target position. This approach can be viewed as a reactive approach since there is no deliberation involved and it can be implemented on lower layers of control. Furthermore, this reactiveness allows us to deal with obstacles on a real-time basis, the only limitation being the time needed to detect and identify objects as obstacles or goals.

While manipulating objects, both proprioceptive data (motor current consumption, motor angular position) and exteroceptive data (color of the object, number of corners detected on the object) are recorded. Both of these types of data are combined to form a vector $\vec{X} = [\vec{S}_P, \vec{S}_E]$. The resulting vector is statistically analyzed in order to extract important parameters that distinctively describe the object being manipulated. For example, the average power consumption of the motors during the manipulation phase will differ depending on an object's weight. This data is then clustered using the standard k-means algorithm [23] and the resulting clusters are labeled.

Prior to clustering, each element of a data vector is normalized using

$$x'_i = \frac{x_i - \overline{x}_i}{\sigma_i} \tag{1}$$

where $i = 1, \cdots, L$ and L is the length of a data vector $\vec{X}$. The mean $\overline{x}_i$ and variance σ_i^2 are calculated with respect to the training data using

$$\begin{array}{rcl} \overline{x}_i & = & \frac{1}{N}\sum_{n=1}^{N} x_i^n \\ \sigma_i^2 & = & \frac{1}{N-1}\sum_{n=1}^{N}(x_i^n - \overline{x}_i)^2 \end{array} . \tag{2}$$

where N is the number of data vectors in the training set. This normalization process is necessary since the elements of a data vector typically have magnitudes that differ significantly.

The labeled clusters are then used to train a radial basis function network [24] (a subsystem of the **sensory perception processor** module) for classifying the clusters based on proprioceptive and exteroceptive data. Rather than choosing a subset of data points of the clusters as the centers of basis functions, we use the k-means clustering algorithm (in which the number of centers must be decided in advance) to determine for each cluster a set of centers which more accurately reflects the distribution of the cluster's data points. The appropriate number of center points is determined by the performance of the resulting network on a validation set. In the implemented neural network, we used a Gaussian function as a basis function. Figure 5 shows the topology of the radial basis function network used for data classification.

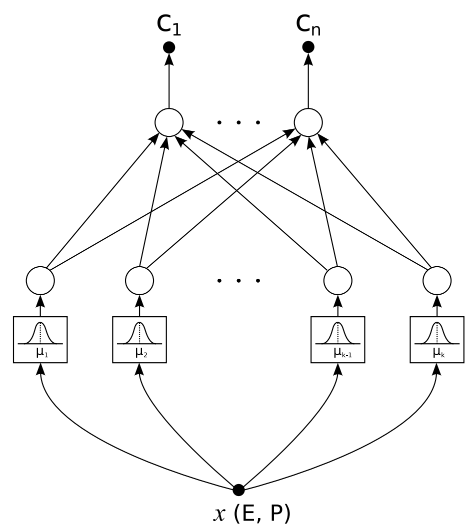

Fig. 5. Radial basis function network

D. Experimental Setup

The robot used for testing our system is one which has been developed in our group, and it is based on the design of the ARAMIES robot [25]. Our robot is a fully functional ambulating robot that is robust and kinematically flexible. It is equipped with various sensors that enable it to perceive both proprioceptive and exteroceptive signals. On each of the robot's legs, there are 6 D.C. motors, 6 pressure sensors, and an infrared sensor. For our experiment, the camera of the robot was used as a source of exteroceptive data, and the average motor current consumption of each motor was used as a source of proprioceptive data.

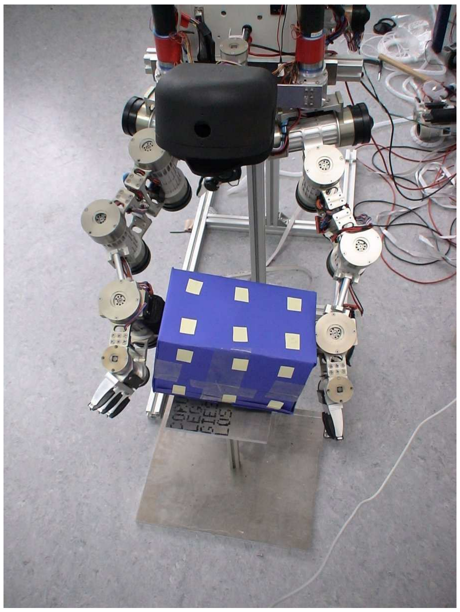

Fig. 6. The robot manipulating an object

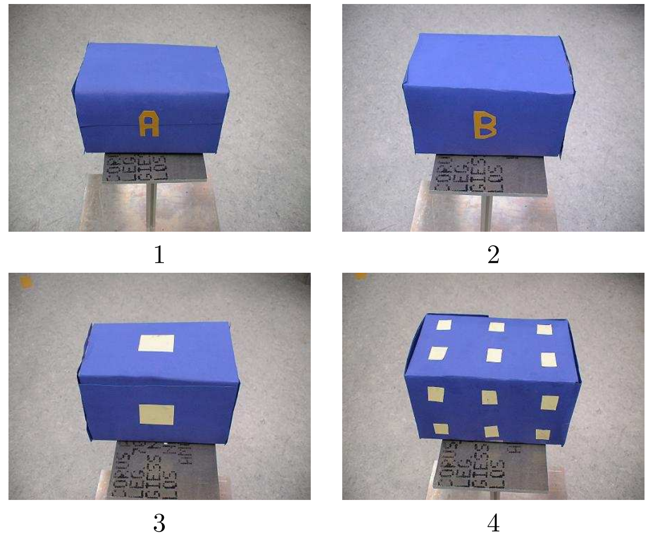

Fig. 7. Objects used in the sensorimotor-coordination experiment

Obj.	$\overline{N_c}$	$\overline{A_c}$	$\overline{\sum I}$	σN_c	σA_c	$\sigma(\sum I)$
1	1	4812.2	4688.68	0	48.22	217.93
2	1	4925.77	5242.52	0	61.53	159.14
3	2	3134.15	4670.66	0	39.5	181.27
4	6.96	953.1	4916.75	0.21	10.4	319.41

TABLE I

The average and standard deviation of features over the whole training set

In the experiment we performed, the robot's body is fixed and it uses its forelegs to manipulate the different objects shown in Figure 7. The objects have differing weights and visual features. Two of the objects have the same visual features, and cannot be distinguished from each other using only visual information; these objects are marked as "A" and "B" in Figure 7. The faces on which the letters are written are placed away from the camera of the robot so that the two objects appear indistinguishable to the robot.

In the training session, five manipulation acts were performed on each of the objects. For a single manipulation act, we took a series of images from which we calculated the average number of contours extracted and the average area of the extracted contours. Furthermore, we calculated the total current consumption average for the motors on both of the robot's forelegs.

E. Results

E.1 Repeatability of Features

Table I shows, for each of the objects in Figure 7, the average and standard deviation of the number of detected contours N_c, the area (number of pixels) of the detected contours A_c, and the total current consumption I (in mA) of both of the robot's forelegs over all training sessions. This data is an indirect measure of the repeatability of a particular feature's measurements. A measurement for a feature is repeatable if the variance of the measurement over a given sample of measurements is small enough that the overlap of measurements resulting from different objects is minimal. One can easily see that the number of contours detected is the most stable feature in this experiment. For getting the number of contours, we used a detector which is robust against noise and changes in lighting conditions. The average current consumption of both legs shows the highest variance in relation to the other features since the end effectors of the forelegs do not grab the object at the same point for each training session. This causes the object's center of gravity to shift with respect to the end effector, and thus a variation in the average current consumption is observed.

E.2 Recognition Rates

We tested the system's ability to recognize the objects it was trained for. The system was tested in three different scenarios. In the first scenario, the system was permitted to use both exteroceptive and proprioceptive data to recognize objects. In this case, the recognition rate was the highest, yielding only one misclassification in 20 trials. In the case where the system was allowed to use only exteroceptive data, there were 7 misclassifications in 20 trials. This poorer performance is explained by the fact that two of the objects have the same visual features. In contrast to these results, when only proprioceptive data was used, there were only 3 misclassifications in 20 trials because the weights of each object were unique. An interesting point is that the system was able to correctly classify objects "A" and "B" in this case, which would have caused problems when using only exteroceptive data.

IV. Conclusions and Outlook

The "Embodied AI" approach assumes that embodied representations play important roles for higher level spatial cognition. The conducted experiments give us first hints on how to test this hypotheses. In the first case study we have implemented and evaluated an approach to automatically classify a spatial environment based solely on the proprioceptive data of an 8 legged, 24 DOF walking robot. The system was able to distinguish beteen three classes of subtrates. Sand and gravel were grouped into a single class due to the fact that for the SCORPION robot sand and gravel have quite the same properties.

In the second case study, an embodied recognition system has been presented which learns to recognize objects by interacting with them and using proprioceptive and exteroceptive data. We have shown that a learning system trained with multimodal sensory information can recognize objects by using only partially available (i.e. only exteroceptive, or only proprioceptive) sensory data. The direct byproduct of such systems is a robust system which continues to operate in the absence of either the proprioceptive or the exteroceptive data. Our preliminary results demonstrate that this method can be effectively used in a robotic system which learns from experience about its environment.

We plan to extend the system by increasing the number of proprioceptive and exteroceptive object-features extracted from the environment, and improving their stability. For example, we may use local features such as Scale Invariant Feature Transform (SIFT) features [26] that describe objects distinctively and which are stable against translation, rotation, scaling and different lighting conditions. Additionally, we will develop a variety of new manipulation acts and work on methods of removing visual ambiguities through manipulation acts. We will apply information reduction techniques such as Principal Component Analysis (PCA) to determine the significance of each individual exteroceptive or proprioceptive sensory input, identifying those which play an important role in recognition for a given manipulation act.

Acknowledgments

The authors would like to acknowledge Fred Laberge for his inputs while developing the learning architecture proposed in this work. Part of this work is supported by the SFB/TR8 DFG project, which is duly acknowledged.

References

[1] K. E. Adolph and S. A. Joh, "Motor developments: How infants get into the act," in *Infant Development*, A. Slater and M. Lewis, Eds. Second Oxford University Press, http://www.psych.nyu.edu/adolph/PDFs/SlaterChapter.pdf, 2006, to appear.

[2] N. Sharkey and T. Zeimke, "The ins and outs of embodied cognition," in *Hybrid Neural Systems*, S. Wermter and R. Sun, Eds. Springer, 1999.

[3] S. Nolfi and D. Parisi, "Exploiting the power of sensory-motor coordination.," in *Advances in Artificial Life, Proceedings of Firth European Conference on Artificial Life (ECAL)*, F. Mondada D. Floreano, J-D. Nicoud, Ed. 1999, pp. 173–182, Springer Verlag.

[4] R. Pfeifer and C. Scheier, *Understanding Intelligence.*, MIT Press, Massachusetts, London, 1999.

[5] G. E. Edelman, *Neural Darwinisim: The Theory of neuronal Group Selection.*, Basic Books, New York, 1987.

[6] R. D. Beer, "Towards the evolution of a dynamical neural networks for minimally cognitive behavior.," in *Proceeding of the 4th International Conference on Simulation of Adaptive Behavior*, 1996, pp. 421–429.

[7] S. Nolfi, "Adaptation as a more powerful tool than decomposition and integration," in *Proceeding of the Workshop on Evolutionary Computing and Machine Learning*, 1996.

[8] M. Takác, "Categorization by sensory-motor interaction in artificial agents," in *Proceedings of the 7th International Conference on Cognitive Modeling*, 2006.

[9] M. Laungerella and R. Pfeifer, "Robots as cognitive tools:Information-theoretic analysis of sensory-motor data," in *Proceeding of the first Int. Conf. on Humanoid Robots*, 2001, pp. 245–252.

[10] F. Kirchner, D. Spenneberg, and R. Linnemann, "A biologically inspired approach towards robust real world locomotion in an 8-legged robot," in *Neurotechnology for Biomimetic Robots*, J. Ayers, J. Davis, and A. Rudolph, Eds. MIT-Press, Cambridge, MA, USA, 2002.

[11] B. Fritzke, "Growing cell structures - a self-organizing network for unsupervised and supervised learning," *Neural networks*, vol. 7, no. 9, pp. 1441–1460, 1994.

[12] D. Spenneberg and F. Kirchner, "Learning spatial categories on the basis of proprioceptive data," in *Proceedings of the 3rd International Symposium on Adaptive Motion in Animals and Machines*, September 2005.

[13] K. W. Spence, *Behavior theory and conditioning*, Yale University Press, New Haven, 1956.

[14] R. A. Wise, "The dopamine synapse and the notion of pleasure centers in the brain," *TINS*, pp. 91–94, April 1980.

[15] R. J. Beninger, "The role of dopamine in locomotor activity and learning," *Brain Research Reviews*, vol. 6, pp. 173–196, 1983.

[16] R. A. Depue and W. G. Iacono, "Neurobehavioral aspects of affective disorders," *Annual Review of Psychology*, vol. 40, pp. 457–492, 1989.

[17] G. J. Mogenson, D. L. Jones, and C. Y. Yim, "From motivation to action: Functional interface between the limbic system and the motor system.," *Progress in Neurobiology*, vol. 14, pp. 69–97, 1980.

[18] G. J. Mogenson and A. G. Phillips, "Motivation: A psychological construct in search of a physiological substrate.," in *Progress in psychobiology and physiological psychology*, J.M. Sprague and A.N. Epstein, Eds., vol. 6, pp. 189–243. Academic Press., New York, 1976.

[19] J. Olds, *Drives and reinforcements: Behavioral studies of hypothalamic functions*, Raven Press, New York, 1977.

[20] M. A. Bozarth, "The mesolimbic dopamine system as a model brain reward system.," in *The mesolimbic dopamine system: From motivation to action*, P. Willner and J. Scheel-Krüger, Eds., pp. 301–330. John Wiley., London, 1991.

[21] R. S. Sutton, "First result with Dyna, an integrated architecture for learning, planning and reacting.," *Neural Networks for Control*, W. Thomas, R. S. Sutton and P. J. Werbos, Eds., Bradford Book., 1992.

[22] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," *IEEE International Conference on Robotics and Automation*, pp. 505–505, 1985.

[23] J. MacQueen, "Some methods for classification and analysis of multivariate observations," in *Proceedings of the Fifth*

Berkeley Symposium on Mathematical Statistics and Probability, 1967, p. 281-297.

[24] C. M. Bishop, *Neural Networks for Pattern Recognition*, Clarendon Press, Oxford, 1995.

[25] J. Hilljegerdes, D. Spenneberg, and F. Kirchner, “The construction of the four legged prototype robot aramies,” in *Proc. of CLAWAR 2005*, London, UK, September 2005.

[26] D. Lowe, “Distinctive image features from scale invariant features,” *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.