A discrete computational model of sensorimotor contingencies for object perception and control of behavior

Alexander Maye and Andreas K. Engel

Abstract— **According to Sensorimotor Contingency Theory (SCT), visual awareness in humans emerges from the specific properties of the relation between the actions an agent performs on an object and the resulting changes in sensory stimulation. The main consequence of this approach is that perception is based not only on information coming from the sensory system, but requires knowledge about the actions that caused this input. For the development of autonomous artificial agents the conclusion is that consideration of the actions, that cause changes in sensory measurements, could result in a more human-like performance in object recognition and manipulation than ever more sophisticated analyses of the sensory signal in isolation, an approach that has not been fully explored yet.**

We present the first results of a modeling study elucidating computational mechanisms implied by adopting SCT for robot control, and demonstrate the concept in two artificial agents. The model is given in abstract, probabilistic terms that lead to straightforward implementations on a computer, but also allow for a neurophysiological grounding. After demonstrating the emergence of object-following behavior in a computer simulation of the model, we present results on object perception in a real robot controlled by the model. We show how the model accounts for aspects of the robot's embodiment, and discuss the role of memory, behavior, and value systems with respect to SCT as a cognitive control architecture.

I. INTRODUCTION

Sensorimotor approaches to explain cognitive functions in humans and animals confer actions an integral function for perception. Instead of considering action and perception as interacting, yet separate processes, in this framework they are two constituents of a process leading to the conscious perception and awareness of the environment as well as the own body. At the core of Sensorimotor Contingency Theory (SCT [1]) is the idea that the relationship between actions and the ensuing changes in sensory stimulation, so-called Sensori-Motor Contingencies (SMCs), characterize the object under inspection, and hence, active exploration is required for object recognition. SCT even more radically re-interprets the role of actions for perception by stipulating that exercising SMCs already constitutes the conscious perception, or awareness, of the respective object. The most important consequence is that the brain is not to construct an internal description of the environment surrounding the agent, a re-presentation, but that the environment itself serves as such a description that can be operated with by acting. In vision, for example, a consequence is that the retinal image alone is not sufficient for building a percept, but that the interplay between the saccades of the eye and the resulting changes of the retinal image is what we experience as 'seeing' an object. This nicely explains experimental data on the perception of ambiguous line drawings, in which a correlation between the bi-stable percept and the pattern of gaze movements is observed [2].

Computational models that faithfully implement SCT are scarce. It is important to note that, according to SCT, actions do not only improve perception (e.g., like in the field of active vision), or are properties of objects (e.g., [3]), but are constitutive for object perception. Successful implementations exhibit desirable properties like structuring information in sensory input [4], robustness against accidental changes in sensor morphology [5], or anticipation [6]. The strongest effect, however, that adopting an SCT stance for developing robot control architectures can have, is probably that this framework does not require object recognition and representation in the classical sense (e.g., [7], [8], [9]).

We developed, therefore, a computational model of visual perception in which actions are an integral component of the perceptual process. The basic idea of the model is to consider SMCs, i.e. the law-like relations between actions and contingent changes in sensory stimulation, as multi-step, action-conditional probabilities of future sensory observations. A set of n-th order Markov models is used as a formal description of these probabilities. In this respect the approach resembles the application of Markov decision processes in reinforcement learning. The model does not depend on the definition of a reward, though, and is not employed for decision making. It is discrete in time, action and sensor space.

The two remaining elements of the control architecture we introduce here are a grouping mechanism, that uses temporal, proprioceptive, or value information to cluster object-related SMCs, and an action selection mechanism. Achieving mastery of SMCs, i.e. the process of deploying the acquired knowledge about the relations between actions and resulting changes in sensory stimulation, requires the definition of a strategy, i.e. a method for choosing an action based on learned SMCs. We employ a stochastic policy here.

To study the learning and exercising of SMCs, and the resulting behavior, we developed a computer simulation of an artificial agent that is completely controlled by the computational model of SCT. The agent can change its orientation in a circular arena, and can measure distance in the direction of its heading. In a second study we investigated the properties of the model in a real robot. Using the LEGO

This work was supported by the Sino-German research training group CINACS, DFG GRK 1247/1, www.cinacs.org, and by the EU project 'eSMCs', IST-270212.

The authors are with the Dept. of Neurophysiology and Pathophysiology at University Medical Center Hamburg-Eppendorf, Martinistr. 52, 20246 Hamburg – a.maye,ak.engel@uke.de

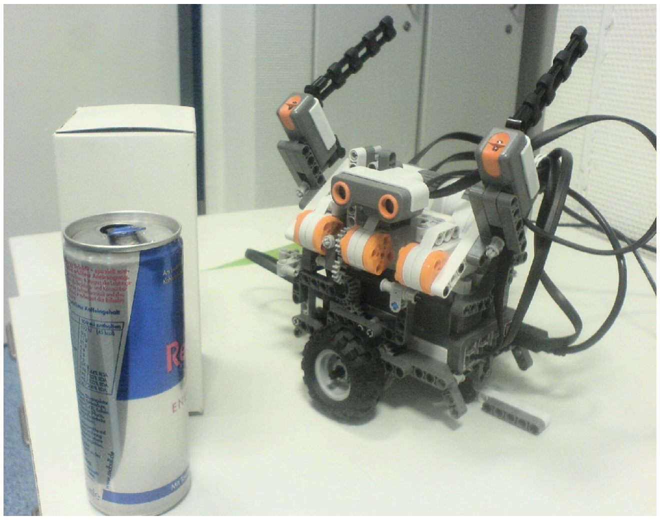

Fig. 1. Lego Mindstorms™ robot for pushing objects to the left or right, and the two objects used in this study (can in the front, box in the back).

Mindstorms™ toolkit, we assembled a mobile robot that can move along a line, and is equipped with an ultrasound distance sensor (Fig. 1). It can use its two independent arms together with its locomotion to move different objects into different directions. The robot's task is to learn to move all cans to the right and all boxes to the left. We show how the robot distinguishes the two object classes by activating different SMCs. This provides reliable object recognition in the face of considerable uncertainty caused by the manual placement of the object in front of the robot by a human, and the significant mechanical instability of the robot.

Conventional approaches to robot control architectures for sorting objects employ a processing pipeline that starts with the acquisition of an image and the extraction of features, e.g. using Scale-Invariant Feature Transform (SIFT [10]). In the next step the position and orientation of objects are normalized in order to facilitate the object recognition in the last step. Then the robot has all the information necessary to select and grasp the object, and move it to the respective position. In sorting 4 household objects, for example, this method can yield an accuracy of around 90% [11]. With the presentation of our approach here we do not primarily intend to compare its performance to that of existing control architectures. This would require to benchmark algorithms using comparable hardware and sensory equipment. Rather we would like to introduce an alternative concept that does not require to reconstruct object type and positioning from static sensory input, and is deemed, therefore, to have advantages with respect to learning efficiency, generalization, and robustness.

II. MODEL

A. Sensorimotor Contingencies

SMCs capture the structure of changes in the sensory input for an action caused by a motor pattern. The theory [1] is not given as a formal description, and indeed only few attempts have been made to derive a computational model of SMCs. In [7] Hebbian learning is employed to increase the synaptic coupling between motor and sensory neuron populations, if they are frequently co-activated. Stability of visual features is used in [9] to adjust these connections. Whereas both approaches employ instantaneous sensorimotor coupling, we decided to develop a model that provides for longer sequences of action-observation couplets.

In our history-based approach a fixed length record of past actions and observations determines the system's current state. At time t the agent makes sensory observation $s(t)$ resulting from action $a(t-1)$. Together with past action-observation pairs, this information constitutes the current state or context $c^h(t) = s(t), \dots s(t-h), a(t-1), \dots a(t-h)$. Based on the current context the agent chooses the next action $a(t)$, leading to observation $s(t+1)$ in the next iteration.

Learning SMCs corresponds to determining the conditional probability of making a sensory observation $s(t+1)$ given an action $a(t)$ and a context $c^h(t)$, and we can write this observation probability as an h-th order Markov model $p^h(s(t+1)|a(t), c^h(t))$.

Different objects will generate different probability distributions, which can partially overlap. The size of the overlap depends on the history length h. Short history lengths will produce distributions that have a large overlap for different objects, and they describe the general effect of actions on the sensor readings independently of the object under consideration. These probabilities model the so-called modality-related SMCs that, according to SCT, give rise to the different experiences of 'seeing' vs. 'hearing' vs. 'touching' etc. With increasing history length, p^h becomes more and more object-specific, corresponding to object-related SMCs in terms of SCT that, when activated, constitute the state of being aware of that object. It should be emphasized that there is no 1:1 mapping between action-observation sequences with large conditional probabilities and objects, i.e. a large p^h_o as such is not indicative for the presence of object o. Rather, several p^h_x become transiently large over time due to the various contexts encountered, and the set $\mathcal{O}$ with $x \in \mathcal{O}$ is the essence of perceiving object o.

This model of SMCs can be implemented by computing histograms of sensory observations $s(t+1)$ for each context, given past action-observation couplets. The probabilty p^h can then be computed from

$$p^h(s(t+1)|a(t), c^h(t)) = \frac{N(s(t+1), a(t), c^h(t))}{N(a(t), c^h(t))} \quad (1)$$

with $N(k)$ being the number of encountering context k. An example for computing these histograms will be given for the simulated agent in section III-A.2. Note that the histogram has to be stored only for contexts that have been actually met, and not all possible contexts. The memory size of this model, therefore, depends on the complexity of the environment, which is a very desirable property. For the studies described in this article we used simultaneously three Markov models of history lengths $h \in [1, 2, 3]$.

B. Value System

A value system allows an agent to select useful actions and avoid futile or even dangerous behaviors, depending on its intentional state and the current context. Reinforcement learning methods are computational models for the development of goal-directed behavior when the agent receives some information about the value of its actions. Value information typically takes the form of reward and punishment, and can be generated internally by the agent, or received from the environment. An interesting approach for generating value information internally is to learn invariances of the sensory input [12]. The output of a set of feature detectors determines the agent's internal state. Using reinforcement learning, the agent finds a set of actions (e.g., camera movements), that minimally change its internal state [9]. This set of actions is representative for the stimulus, and provides content to the sensory state. From the enactive viewpoint [13], the existence of internally generated rewards is questionable, however. In our model we prefer an external reward scheme, therefore, in which rewards and punishment are generated in the environment, and that evaluates the behavior of the agent: When the correct arm is lowered, or the object in front of the robot is pushed to the correct side, the agent receives a reward. Lowering an arm with an object underneath causes a penalty. For all other behaviors (i.e. most of the time) the robot receives no or neutral feedback, which is why conventional reinforcement learning methods are not appropriate for learning SMCs.

The value system employed by the robot is similar to SMCs in that the agent observes conditional reward or penalty probabilities. This is done by counting the number of rewards received in a given context, i.e.

$$p^h_{reward}(s(t+1)|a(t), c^h(t)) = \frac{N_{reward}(s(t+1), a(t), c^h(t))}{N(a(t), c^h(t))}. \quad (2)$$

Another histogram is used for computing the probability of punishment $p_{penalty}$.

Assigning reward probabilities to SMCs can be regarded as a method for structuring the sensorimotor knowledge that the agent acquires. It establishes relations between SMCs that have a similar ecological value. The sensorimotor knowledge could be further structured by considering temporal information, e.g. by grouping SMCs that have been observed in the same time interval and therefore are probably related, or proprioceptive information about the energy consumption of movements to minimize costly actions. In this study we only use reward information from the environment to encourage certain behaviors of the agent, but structuring SMCs in a more complex way using information mentioned before could be used to generate behavior that is more adaptive.

C. Action Selection

In the terminology of SCT, behavior results when the agent actively exercises its mastery of SMCs. However, SCT does not suggest a strategy for choosing an SMC to activate. This is why our model uses the value system described above, which structures and weights SMCs according to their expected reward. The decision schema (Fig. 2) constitutes a simple switch from exploration to exploitation: If there are actions that have never been tried before in the current context, one of these actions will be chosen. Otherwise, an action that most likely will yield a reward is chosen. If there are only actions for which only no or negative rewards (i.e. penalties) have been observed in the past, these actions are assigned probabilities according to their least negative effect, and one of them is chosen according to this probability distribution.

We give a formal description of the action selection schema in the following. The number that an action a has been chosen in the current context is $N(a, c^h(t))$. The set of actions that has never been executed before in the current context then is $\cup_h\{a|N(a, c^h(t)) = 0\}$, and one of these actions is chosen with equal probability. If this set is empty, the reward probabilities for each possible action a are computed as $\max_h \left(\sum_{s(t+1)\in S} p^h_{reward}(s(t+1)|a, c^h(t))\right)$, and the action with the highest reward probability is chosen. If several actions would yield identical rewards, one of them is chosen randomly. Finally, if all reward probabilities are zero, i.e. this context has never yielded a reward, actions are chosen according to $\max_h \left(1 - \sum_{s(t+1)\in S} p^h_{penalty}(s(t+1)|a, c^h(t))\right)$. This ensures that actions incurring a definite punishment are not considered, and the remaining actions are chosen with respect to the probability of having a negative effect.

III. Results

A. Simulated Agent

1) Experimental Setup: We used a computer simulation to demonstrate that the presented model can generate useful behavior in an artificial agent. The agent is located in the center of a circular arena and measures the distance in the direction of its heading (Fig. 3a). During a single action it can rotate approximately $5°$ or $15°$ to the left or to the right. The arena is equipped with two simulated objects,

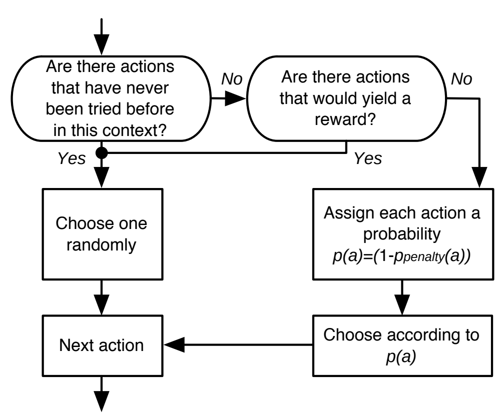

Fig. 2. Schematic of the action selection algorithm based on rewards observed in the past.

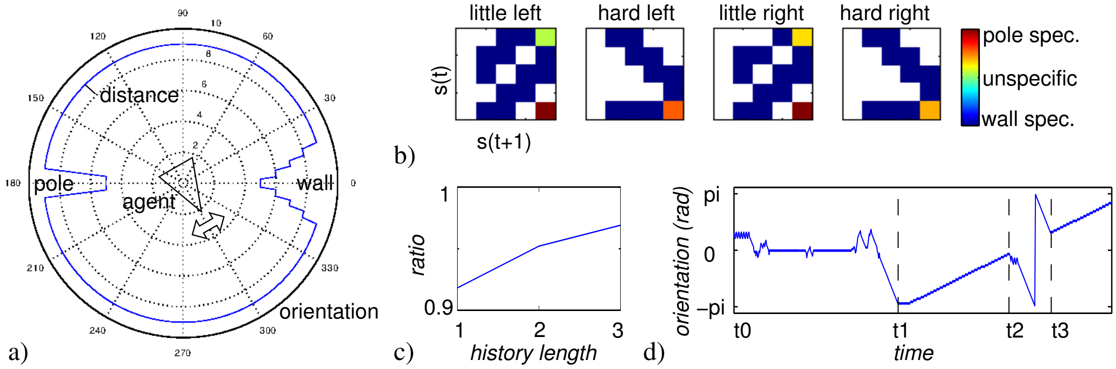

Fig. 3. **a)** Computer simulation of an agent that can change its orientation and measure distance. **b)** History length 1 SMCs (p^1) emerging from learning after 10.000 iterations in an arena containing a pole and a wall. The plots show the conditional probabilities of observing distance sensor reading $s(t+1)$ after executing the action given at the top of each panel, when the current distance reading is $s(t)$. Action-observation tuples (SMCs) in red/blue are exclusively observed when focusing the pole/wall, other colors are less specific, white patches represent tuples that have never been observed. **c)** Ratio of SMCs that are specific for one of the two objects for each of the three history lengths used. The ratio of object-specific SMCs increases with history length. **d)** Plot of the agent's orientation over time in an experiment showing object-following behavior. After associating a reward with the SMCs of the pole for some time, the agent is initialized to a random orientation at t_0. It starts moving randomly, eventually getting sight of the pole located at orientation $-\pi$ at t_1. Then the pole starts moving counter-clockwise, and the agent maintains focus. At t_2 the pole is in front of the wall (having identical distance), and the agent loses track. However, it quickly reorients and continues following at t_3.

a pole, giving a binary distance response only in a very narrow orientation range ($\approx 10°$), and a wall with a graded response in a range of approximately $40°$. For simplicity we used 4 discrete actions and 5 discrete sensor readings for the distance. The agent observed conditional probabilities $p^h, h \in [1, 2, 3]$ in an exploration phase, choosing actions randomly. In the following learning phase, the agent received an external reward whenever it was oriented towards one of the objects.

2) Sensorimotor Contingencies: For the computation of conditional observation probabilities according to eq. (1), the agent maintains a set of histograms of observed contexts. For example, if at time t the agent is heading slightly to the right of the pole, its sensor measures the distance to the background (defined to yield a reading of 1), so $s(t) = 1$. If it would move a little left ($a(t) = 1$), this might not be enough to get sight of the pole, hence $s(t+1) = 1$, and $N(1, 1, 1)$ would be increased by 1. If it would move hard left ($a(t) = 2$), however, it would now sense the distance to the pole ($s(t+1) = 5$), and $N(1, 2, 5)$ would be increased instead. By the same token counts $N(k)$ get increased for any other region of the arena that the agent explores, including the range when it faces the wall. Therefore the agent does not know the object identities when it samples the probabilities during the exploration phase. Only by feedback from its environment it learns to associate reward information with some SMCs, which it then can use to follow the rewarded object.

Fig. 3b shows an overlay of the two sets of first-level history SMCs (p^1) for the two objects. Obviously both sets comprise object-related SMCs as well as SMCs that are shared between the two sets. Fig. 3c shows that the ratio of SMCs that are unique for each of the two objects increases with history length h, confirming the expected relation between probabilities conditional on long histories and object-related SMCs.

3) Object-Following Behavior: The reward schema can result in an object-following behavior of the agent. Switching the action selection strategy from choosing actions randomly to one that tries to maximize p^3, p^2 or p^1 in that order, made the agent follow the formerly rewarded object when this was moved along the perimeter of the arena. As Fig. 3d shows, the agent can maintain this fixation mostly even when the two objects overlap. Orientation is temporarily lost when the objects perfectly overlap, giving the same distance information, but is quickly re-established when they can be disambiguated again.

B. Embodied Agent

1) Experimental Setup: From the Lego Mindstorms™ robotic toolkit we assembled a 2-armed robot that can move along a line (Fig. 1). The discrete action space comprises 4 locomotion actions (little/hard left/right), and 4 arm actions (left/right arm up/down). An ultrasound sensor continuously measures distance to object surfaces in front of the robot during movements. The distance profile recorded during a movement is classified into one of 5 sensor values[1], so that $s(t+1) \in [1, 5]$ for each action $a(t)$.

The robot's drive comprises several gear wheels, causing a noticeable slip. In addition, the slip depends on the action sequence, with sequences that change the movement direction having a much larger slip than sequences moving the robot to the same direction. Together with the general, mechanical instability of the toolkit parts, the robot features a

[1]The 5 conditions are: all raw distance measurements are infinity (nothing in sight), only the first few measurements are infinity (getting sight of something), only the last few measurements are infinity (losing sight), the variance of measurements is larger than 0.2 (uneven surface), the variance is less than 0.2 (even surface). Note that this classification alone does not allow distinguishing the two objects, as the can could yield the value for a flat surface when looked at in the center, and the box could yield the value for an uneven surface when it was not exactly aligned with the direction of movement of the robot.

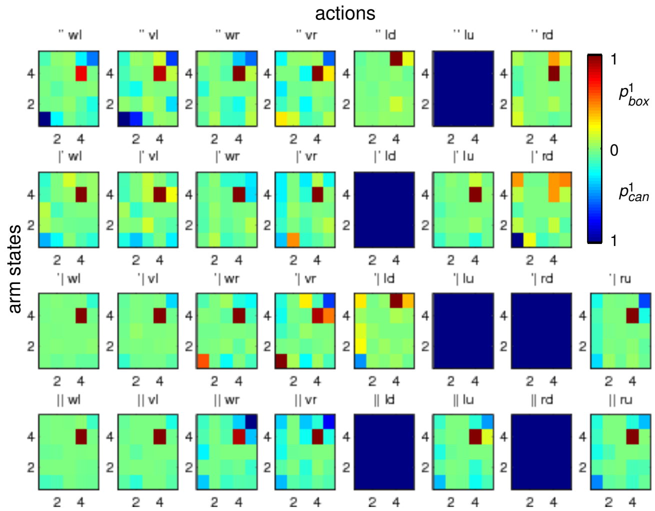

Fig. 4. SMCs of a can (cold colors) and a box (warm colors) for history length 1 (p_o^1). Columns correspond to 8 possible actions, rows to 4 possible states of the 2 arms. Each plot gives the probability of observing s(t+1) (x-axis) given s(t) (y-axis). Note how the flat face of the box generates a high probability of repeatedly observing sensor value 4, and how the can generates several different contexts. Empty plots correspond to impossible actions.

degree of unreliability in its locomotion that would be very undesirable from a robotic viewpoint. This is an excellent setup, however, for assessing how SCT accounts for the particular properties of the agent's body in particular, and for studying the implications of SCT in the control of real agents in general.

2) Object Perception: By moving left and right, the robot could explore two objects, a can and a cardboard box (Fig. 1), that were placed in front of it successively. For display purposes we separated the SMCs for each object in two sets $p_o^h(o \in \{can, box\}, h \in [1, 2, 3])$ of SMCs in this experiment (Fig. 4), but the robot always evaluated the union $p_{can}^h \cup p_{box}^h$. It is interesting to note that the strong symmetry between SMCs for inverse actions (e.g., little left and right), that was observed in the simulation, is broken in the robot. We interpret this as the result of the process of shaping SMCs by the robot's embodiment.

After the learning phase, the two objects were presented to the robot at random. The conditional probabilities for the new action-observation sequences were integrated separately for the two sets of SMCs over time to build up and visualize the belief about the identity of the object perceived by the robot. As can be seen in Fig. 5, the robot quickly recognized the object. An interesting observation made is that the belief is not developing monotonically. It depends on the object part the robot is currently inspecting. At the object borders, where distance measures jump equally for both objects, separation between the respective probabilities levels out. The separation increases, though, when the robot inspects the central parts, where both objects show the strongest differences in the distance profile.

3) Goal-Directed Behavior: To demonstrate the emergence of useful behaviors in a robot controlled by our model of SMCs, we applied an operant conditioning paradigm. The robot received a reward for lowering the left arm or moving to the right with the left arm lowered, if it was exploring a can, and vice versa for the box. It received a punishment when it tried to lower an arm with an object underneath. This reward scheme effectively conditions the robot to push the can and the box to opposite sides, and to avoid obstacles when lowering an arm.

In reality rewards would be given by an external observer to the robot. Since the reward conditions are simple conjunctions of the robot's state and actions, which are both known to the control program, we implemented the reward mechanism as a sub-function in the program that simulates the stereotyped behavior of an observer. Punishments are the negative value assigned to proprioceptive feedback when lowering an arm is impeded. Since the robot is not equipped with touch or force sensors, we simulated this feedback by comparing the control commands for the arm motors with their actual position and checking if they agree.

We consider a behavior that tries to maximize reward as goal-directed. During the learning phase, the agent playfully pushes the object for about 5 times in sequence (see Fig. 6). After playing for about 3.000 epochs with the cardboard box and 11.000 iterations with the can, it has learned that it can maximize reward by deliberately pushing the object for more episodes. The maximum number of consecutive rewards increases drastically in this exploitation phase. Since the can features more complex SMCs than the box, a longer exploration phase emerges.

IV. CONCLUSIONS AND FUTURE WORKS

A. Conclusions

In this article we introduce the notion of considering SMCs as conditional probabilities over a joint action and observation space, and of using a set of Markov models to capture different context sizes, given by different history lengths of action-observation couplets. Different objects generate different probability distributions in different agents, hence SMCs are object- as well as agent-specific. Our model

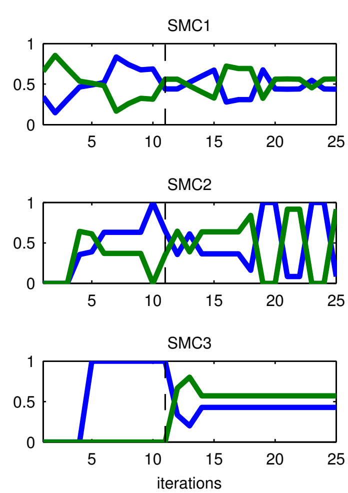

Fig. 5. To test the capabilities for object recognition by exercising its acquired SMCs, we present the robot with one of the two learned objects and record the conditional probabilities for each object over time (blue: can, green: box). After 10 iterations the objects are exchanged (dashed line). Note how increasing history length (SMC1→SMC2→SMC3) leads to a more and more sustained larger probability for the correct object identity, hence showing evidence for object-related SMCs.

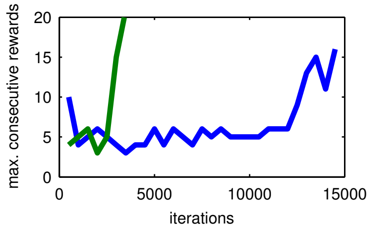

Fig. 6. Maximum number of consecutive episodes of pushing the current object (blue: can, green: box) in a specific direction.

suggests that the two types of SMCs, modality-related and object-related SMCs, may be captured by different history lengths. In combination with a value system and a controller that evaluates the set of learned SMCs, the agent can show goal-directed behavior. The robot reliably recognizes an object despite the significant mechanical instabilities of its body, which cause a large variation in the data collected from an object. The behavior is also robust under dynamic changes in the environment, like small variations in object position after arm contacts, or replacing an object under inspection with another by the experimenter. The results presented here are encouraging and show that SCT is not only an attempt to advance our understanding of cognitive capabilities in biological agents, but also provides a framework for the development of alternative robot control architectures.

B. Future Works

In order to compare the performance of our approach to existing systems, a statistical evaluation of the sorting accuracy would be necessary. In particular we will be interested in quantifying the reliability against variability in positioning, and in exploring the generalization capabilities. Likewise the scalability to more object classes are of major interest. Employing Markov models with increasing history lengths should allow to separate more object classes. In theory the number of distinguishable objects should grow exponentially with history length, but since several SMCs are learned for each object class the actual capacity will grow slower.

For a large-scale applications the simple switch to the exploitation of acquired knowledge only after all possible actions in all observed contexts have been explored will not work due to the possible combinatorial explosion between large context and action spaces. We do not consider this a major problem, though. First, effective methods to switch between exploration and exploitation exist (e.g., based on intrinsic motivation [14]), and can be combined with our model. And second, our approach does not require a complete exploration of all SMCs for all objects. Rather the accuracy of the robot's behavior gradually increases with the number of learned SMCs, which allows for a continuous improvement of performance during task execution, provided some feedback is given to the robot. For example, of the $5^4 \times 8^3$ possible SMCs with history length 3 in our study, the robot actually observed only $\approx 0.05\%$, which is sufficient to distinguish the two objects.

The model we present here is a first step towards a computational realization of the concepts proposed by SCT, and main directions of further development are evident. Clearly, discrete actions and sensor readings are not biologically realistic, so the model should be based on continuous actions and sensor spaces instead. Extending the model in this respect seems feasible. SCT has been conceived to explain human perception, and the model may have a value in understanding the function of perceptual processes in the human brain. To ground the current abstract computational model in neurophysiological processes, a neuronal mechanism for learning conditional probabilities has to be employed. While we use multidimensional histograms here, neurobiologically plausible implementations are evident (e.g., Radial Basis Function Networks). For robotic applications, implementations of the model need not to be neurobiologically plausible. These applications can be fully optimized for the computer hardware they will be running on, and the general formulation of the model introduced here provides a large potential for effective implementations.

V. ACKNOWLEDGMENTS

The authors gratefully acknowledge the valuable comments of Jianwei Zhang on the manuscript.

REFERENCES

[1] J.K. O'Regan and A. Noë. A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, 24:939–973, 2001.

[2] T.C. Kietzmann, S. Geuter, and P. König. The role of overt visual attention in the process of perceptual formation. In *Rovereto Attention Workshop*, 2009.

[3] T. Asfour, A. Welke, P. Ude, P. Azad, J. Hoeft, and R. Dillmann. Perceiving objects and movements to generate actions on a humanoid robot. In *IEEE International Conference on Robotics and Automation (ICRA'07)*. Workshop: From features to actions - Unifying perspectives in computational and robot vision, April 2007.

[4] M. Lungarella and O. Sporns. Mapping information flow in sensorimotor networks. *PLoS Computational Biology*, 2 e144(10):1301–1312, 2006.

[5] P. Fine, E. Di Paolo, and E. Izquierdo. Adapting to your body. In *Ninth European Conference on Artificial Life*, pages 203–212. Springer, 2007.

[6] H. Hoffmann. Perception through visuomotor anticipation in a mobile robot. *Neural Networks*, 20:22–33, 2007.

[7] S. Bovet and R. Pfeifer. Emergence of delayed reward learning from sensorimotor coordination. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2272 – 2277, August 2005.

[8] M. Maillard, O. Gapenne, L. Hafemeister, and P. Gaussier. Perception as a dynamical sensori-motor attraction basin. In *Advances in Artificial Life (8th European Conference, ECAL)*, pages 37–46, 2005.

[9] Y. Choe, H.-F. Yang, and D.C.-Y. Eng. Autonomous learning of the semantics of internal sensory states based on motor exploration. *International Journal of Humanoid Robotics*, 4:211–243, 2007.

[10] D.G. Lowe. Distinctive image features from scale-invariant key-points. *International Journal of Computer Vision*, 60:91–110, 2004.

[11] A.C. Romea, D. Berenson, S. Srinivasa, and D. Ferguson. Object recognition and full pose registration from a single image for robotic manipulation. In *IEEE International Conference on Robotics and Automation (ICRA '09)*, May 2009.

[12] L. Wiskott, M. Franzius, P. Berkes, and H. Sprekeler. Is slowness a learning principle of the visual system? In *Proc. 39th Annual European Brain and Behaviour Society, Triest, Italy, September 15-19*, pages 14–15, 2007. Special issue of Neural Plasticity, Article ID 23250.

[13] E. Di Paolo, M. Rohde, and H. De Jaegher. *Enaction: Toward a New Paradigm for Cognitive Science*, chapter Horizons for the enactive mind: Values, social interaction and play, pages 33–87. MIT Press, Cambridge, MA, USA, 2010.

[14] F. Oudeyer P.-Y., Kaplan and V. Hafner. Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation*, 11(2):265–286, 2007.