Time Scales of Sensorimotor Contingencies

Alexander Maye and Andreas K. Engel

University Medical Center Hamburg-Eppendorf
Dept. of Neurophysiology and Pathophysiology
Martinistr. 52, D-20246 Hamburg, Germany
{a.maye,ak.engel}@uke.de

Abstract. In Sensorimotor Contingency Theory (SMCT) differences between the perceptual qualities of sensory modalities are explained by the different structure of dependencies between a human's actions and the ensuing changes in sensory stimulation. It distinguishes modality-related Sensory-Motor Contingencies (SMCs), that describe the structure of changes for individual sensory modalities, and object-related SMCs, that capture the multisensory patterns caused by actions directed towards objects. These properties suggest a division of time scales in that modality-related SMCs describe the immediate effect of actions on characteristics of the sensory signal, and object-related SMCs account for sequences of actions and sensory observations. We present a computational model of SMCs that implements this distinction and allows to analyze the properties of the different SMC types. The emergence of perceptual capabilities is demonstrated in a locomotive robot controlled by this model that develops an action-based understanding for the size of its confinement without using any distance sensors.

1 Introduction

One of the main propositions of SMCT [1,2] is that the brain neither builds nor uses internal representations of the world to generate sensory awareness. Instead the theory suggests, that the interrelations between actions and changing stimulation patterns impinging on the sensory organs while exploring the environment constitute its perceptual contents and conscious experience. The idea that action plays a constitutive role in perception has been developed by a school of philosophers taking a pragmatic stance on cognition (see [3] for an overview). In robotics, the construction of reliable, internal world models and their sensible employment for the generation of behavior are major challenges, that to date have not been satisfied. Therefore SMCT should be highly relevant for the development of control architectures for artificial agents with cognitive capabilities. However, surprisingly few approaches implement the core concept of SMCT [4,5,6,7]. Many more approaches are *inspired* by SMCT, but in our view fail to realize the constitutive role of action for perception.

SMCT postulates two types of SMCs, namely modality- and object-related SMCs. Inspired by studying a computational model of SMCs that was proposed by the authors previously [8], we argue here that these two SMC types "live" on

H. Zhang et al. (Eds.): BICS 2012, LNAI 7366, pp. 240–249, 2012.

different time scales. This idea will be further motivated in the next section. In section 3 we present the results of a study in which a robot learned the SMCs defined by the specific properties of its embodiment and the environment, and used them to optimize its behavior. The study demonstrates how the longer temporal context of object-related SMCs can be used as the basis for knowledge about the size of the robot's confinement. To make the case more interesting, the robot does not use sensors that measure distance directly (e.g. infra-red distance sensors) or indirectly (e.g. camera images). In the last section we discuss implications and hypothesize about possible extensions of this approach.

2 Time Scales of SMCs

2.1 Two Types of SMCs

The different sensory organs of the human body measure different physical quantities: luminance at the retina, sound pressure at the ear-drum, force at the tactile skin receptors and in the muscle spindles, accelerations in the vestibular system etc. It comes as no surprise, therefore, that the structure of changes in the sensory signal caused by actions is different for each sensory modality. O'Regan and Noë ([1] p. 941) mention the differences between vision and audition: while eye movements generate changes in the sensory stimulation on the retina, they do not affect the sensory signal coming from the ear. On the other hand, movements of the head change the temporal asynchrony of sounds hitting the ear-drums, which is the basis for sound localization. Of course head movements change the luminance distribution on the retina as well, but the pattern is different from both, the changes caused by eye movements and the changes in temporal asynchrony of sound waves. They conclude that

> ... what does differentiate vision from, say, audition or touch, is the *structure of the rules* governing the sensory changes produced by various motor actions, that is, what we call the *sensorimotor contingencies* governing visual exploration. Because the sensorimotor contingencies within different sensory domains (vision, audition, smell, etc.) are subject to different (in)variance properties, the structure of the rules that govern perception in these different modalities will be different in each modality. (p. 941, emphasis theirs)

This structure of rules is what we call modality-related SMCs. Since they are a direct consequence of the "(in)variance properties" of the physical quantities as such, contingent on the agent's actions, we conjecture that they capture the instantaneous effects of these actions on the patterns of sensory stimulation.

The second type of SMCs that O'Regan and Noë identify are determined by (visual) attributes of objects. When inspecting an object, the signals in the visual domain change in a lawful way that is characteristic for the object at hand. An example is the change of light reflection and color when the object is manipulated, and the fact that we have to turn it around to be able to perceive its backside. They suggest

> ... that the visual quality of shape *is precisely* the set of all potential distortions that the shape undergoes when it is moved relative to us, or when we move relative to it. Although this is an infinite set, the brain can abstract from this set a series of laws, and it is this set of laws which codes shape. (p. 942, emphasis theirs)

The fact that these object-related SMCs require some kind of exploration suggests that they are activated on a longer time scale, and that they establish a form of context. To perceive an object requires to explore its attributes, be they visual or tactile, in a process that extends over longer times than the rather direct modulation of the signals in the different sensory modalities by the actions in this process. This is why we think that the different types of SMCs are associated with different time scales.

2.2 Accounting for Different Time Scales in a Computational Model of SMCs

The idea to consider the relations between types of SMCs and time scales has been inspired by studying a computational model of SMCs that was previously proposed by the authors. The details are given in [8], but the general idea is pictured here again. While the agent executes an action a, sensory observations o are recorded. The model works in discrete time, so after each time step a new action-observation pair $ao(t)$ – an SMC – is available. These action-observation couplets are pushed on a first-in-first-out queue that keeps a memory of recently executed actions and resulting sensory observations (see Fig. 1). An array of Markov models is employed to learn the probability distributions of patterns in this queue, corresponding to the different action-observation contexts the agent experiences. Each Markov model takes a finite history of actions and observations into account[1]. These probability distributions are a model of the SMCs of the respective agent. Since the Markov models with short history lengths disregard the longer context of specific situations that the agent goes through, they capture general regularities of the sensorimotor patterns. In this respect they seem to have similar properties like modality-related SMCs. In contrast, the models that take a longer history of actions and observations into account, represent sensorimotor patterns that depend on the specific context. Exploration of the visual attributes of an object is an example of such a context, and hence Markov models with longer histories are deemed to realize object-related SMCs.

We would like to point out that in these considerations we do not make any assumptions about the number of sensory modalities involved. In particular we do not interpret the term modality-related in a way that suggests this SMC type would capture structural regularities of each sensory modality in isolation. Rather than employing different subsystems for different modalities, our approach is based on the joint probability distributions of all sensory modalities. Therefore differences in the sensorimotor laws between individual modalities

[1] Note that the probability distribution for history length h can be determined from the one for $h+1$ by computing the marginal distribution over all $ao\,(t-(h+1))$.

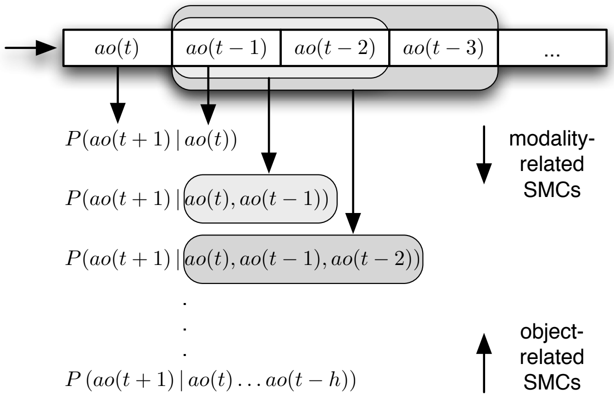

Fig. 1. Schema of the proposed computational model of SMCs. The sequence of actions and sensory observation (ao) is used to train Markov models of different history lengths h. We suggest that object-related SMCs are captured by models with longer context sizes, and modality-related SMCs by models with a short context.

are reflected in the qualitative differences of projections of the joint probabilities to (the dimensions of) individual modalities, and not in physically distinct, modality-specific subsystems. We believe that this agrees well with the very idea of SMCT. Technically, though, it is still possible for an external observer to visualize strictly unimodal SMCs by computing marginal probability distributions.

3 Robot Study

In this section we present a particular implementation of the model described in the previous section. It demonstrates the simultaneous acquisition of modality- and object-related SMCs, how the former control efficient instantaneous behavior, and how the latter lead to meaningful long-term behavior.

3.1 Hardware

The Robotino® robot (Festo Didactic, Esslingen, Germany) is equipped with an omnidirectional (Swedish) wheel drive, a webcam, 9 distance sensors, and a collision detector. For this study locomotion was restricted to forward and backward movements at constant speed. Camera and distance sensors were not used, because in the experimental setting (described below) the robot should rely only on the proximal senses that are described in the following. The collision detector is implemented in the Robotino® by a compressed air tube, that is attached around the circular periphery and registeres pressure changes. It does not yield directional information about the side of the collision. Readings of the instantaneous current consumption of the motors were used to detect whether the robot pushes against an obstacle that had not triggered the bumper. A custom-made accelerometer was attached to the chassis of the robot giving three-dimensional acceleration information.

The model for controlling the robot's behavior uses discrete time steps, and we set the update rate to 1 Hz. Every action (move forward or backward) was executed for 1 second. During this time sensor values were sampled at 10 Hz. At the end of each epoch the raw sensor values were averaged and quantized. This preprocessing step did not serve to recognize or classify the sensory input, as is the case in conventional robot control architectures, but to reduce the complexity of the sensory signals to a level that allows fast learning of SMCs.

3.2 Experimental Setting

The robot was placed on a flat ground between two walls at about 4 times the diameter of the robot apart. Traveling from one wall to the opposite took about 5 seconds. While the wall in the back of the robot triggered the bumper on collisions, the wall in the front had a protrusion that prevented further forward movement, but did not trigger the bumper (Fig. 2). The robot should learn to move in an smooth and energy-efficient way, and to respond appropriately to collisions. Collision avoidance is typically implemented in the lowest levels of conventional control architectures, relying on bumpers or distance sensors. However, in the future we would like to employ our approach in scenarios in which the robot is allowed to push objects around. A reflex-like obstacle-avoidance mechanism could not be used in this case. Instead the robot must have the opportunity to observe the sensory feedback from the bumper, the motors, and the distance sensors to feel if it is pushing an object or bumping into a wall.

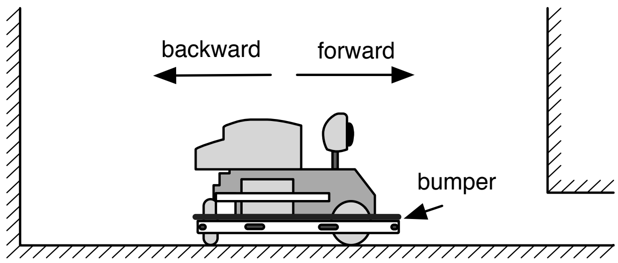

Fig. 2. Profile view of the experimental setup. The robot could freely move forward and backward between the two walls. The shape of the front wall (right) stopped the movement, but did not trigger the bumper.

An analogous setting for humans would be a confinement in the darkness, where the prisoner can only walk forward and backward. After exploring the prison cell for a while, he would know that it takes 5 steps from wall to wall, and while the one in the back feels hard, the one in the front is soft.

3.3 Model

We employed a slightly modified version of the computational model of SMCs described in [8]. Apart from the sensory preprocessing, the value system and the action selection schema had to be adapted to the Robotino's® embodiment.

Value System. Defining a value system is a method "to make an agent do something in the first place" [9]. Here we used only internal values as the experimental setting did not feature any rewards coming from the environment. The features from the bumper ($\in 0, 1$), motors ($\in 0, 1, 2$), and accelerometer ($\in [-2, +2]$) contributed to the internal value in the following way:

$$v_{int} = -bumper - 0.5motor - 0.2 \max_{x,y,z}(|accel|)$$

The internal value was always less than zero, and the robot should try to find a behavior that makes it least "negative".

Action Selection. In addition to the probabilities of action-observation pairs ao conditional on a history of length h of previous action-observation pairs, $P^h\left(ao(t+1)|ao(t)\dots ao(t-h)\right)$, the model keeps track of the average value associated with each context of size h:

$$\bar{v}^h_{int}\left(ao(t+1)|ao(t)\dots ao(t-h)\right) = \frac{\sum v_{int}\left(ao(t+1)|ao(t)\dots ao(t-h)\right)}{N\left(ao(t)\dots ao(t-h)\right)}$$

with $N()$ being the number of times having experienced this context. Using this value probability, a simple schema to select the next action is used: Working down from large context sizes h to smaller ones, for each possible action the average value is found. This requires that the current context, given by the sequence $ao(t)\dots ao(t-h)$ has been experienced before, i.e. $N\left(ao(t)\dots ao(t-h)\right) > 0$. Now the action with the best predicted value gets executed with probability $p = \bar{v}^h_{int} + 1$. Actions that reliably have led to a perfect internal value ($\bar{v}^h_{int} = 0$), for example, would be readily executed again. Actions with a average internal value at or below -1 would never get executed. All intermediate values leave the opportunity to explore alternatives, with a probability depending on the previous experience with this action.

3.4 Results

We can evaluate the behavior of the robot by comparing the components of the value function when it is controlled by the just described SMC-based action selection schema, and when actions are selected randomly. The plots in Fig. 3 show the time course of the internal value components for about 35 minutes run time of the robot. During the first 5 minutes the robot has learned that fewer switches between forward and backward drive lead to fewer acceleration peaks and a reduced energy consumption, and hence improve the internal value. Stability of motor current, acceleration and turning probability between 5 and about 10 minutes runtime characterizes the conclusion of the learning phase

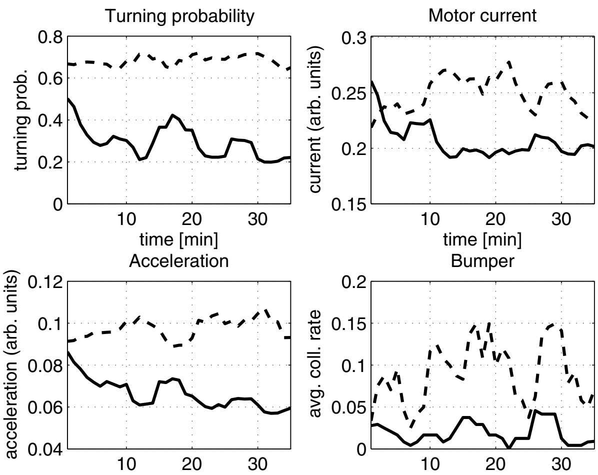

Fig. 3. Time course of the sensory signals in the three modalities that affect the internal value associated with each SMC: acceleration, current consumption, and bumper. The solid line is from a run with the described action selection schema, the dashed line resulted from randomly switching forward and backward drive. Turning probability is not considered for the internal value, but gives an idea how straight the robot moved (low values). All curves are smoothed with a moving average window of 4 minutes.

of modality-related SMC. The robot now knows how to move in an efficient manner, but still tries to improve its internal value and explores different action sequences. This is achieved at around 12 minutes when it discovers the full extent of its confinement, allowing to move with a minimum number of turns. The robot continues to switch occasionally to exploration (e.g. at around 25 minutes) for two reasons. First it might have experienced a hitherto unknown SMC and starts searching for appropriate responses. And second, the turn it makes in front of a wall to avoid the collision still incurs an “unpleasant” sensation of acceleration. Therefore it checks out the only alternative action, resulting in a collision.

One might be interested in the question why the probability of triggering the bumper shows no significant change over time. This has to be analyzed in conjunction with the motor current. The combination of high collision frequency and high currents in the beginning indicates that the robot frequently pushed against the wall for several epochs. It had not yet learned the appropriate behavior upon a collision, which is to immediately drive in the opposite direction. It shows this behavior after learning has progressed, e.g. around 15 and 25 minutes, when it encounters again phases of frequent collision, but without a corresponding increase in current consumption. The other reason for a sustained rate of collisions is the slip between the wheels and the ground. Even though the robot knows after some time the distance between the walls in terms of the number of consecutive steps in one direction, it has no means to determine its absolute position. The only way to do this is to check the wall once in a while to bring its internal model in register with the world.

An impression of the modality-related SMCs that the robot acquired during the experiment can be gained from Table 1. It shows SMCs in human readable form for a history of length 0, i.e. $P(ao(t+1))$. Most movements required only little current, were not associated with acceleration peaks, nor triggered the bumper. This kind of behavior was of good internal value, and hence was exercised frequently. When the movement in the previous epoch was in the opposite direction, the robot incured an acceleration peak, leading to somewhat lower internal values. Collisions with the walls made for a bad experience and were tried to be avoided.

In Table 2 three examples of SMCs if history length 5 are given. The first block shows the most frequently observed SMC of history length 5 ($p = 0.06$), reflecting a peak in acceleration when the movement direction is changed from forward to backward after 3 time steps. The second block shows an SMC of

Table 1. Listing of the 6 most frequently observed SMCs of history length 0. Most of the time going forward or backward caused a low current consumption, no acceleration peaks, and didn't trigger the bumper (first two rows). The robot "felt" the acceleration when it was changing the direction of movement (third and fourth row). Pushing forward against the protrusion increased the current consumption (fifth row), and pushing backward compounded a signal from the bumper (last row).

Probability	Value	Action	Current	AccelX	AccelY	AccelZ	Bumper
0.41	0.00	→	+	0	0	0	
0.36	0.00	←	+	0	0	0	
0.10	-0.20	←	+	-	0	0	
0.06	-0.20	→	+	+	0	0	
0.03	-0.50	→	++	0	0	0	
0.01	-1.50	←	++	0	0	0	∗

Table 2. Selection of SMCs of history length 6. The arrangement of features has been transposed in comparison to Table 1. Action-observation pairs are now shown in rows, and columns represent different time steps (progressing from right to left).

Feature	$ao(t+1)$	$ao(t)$	$ao(t-1)$	$ao(t-2)$	$ao(t-3)$	$ao(t-4)$	$ao(t-5)$
Action	←	←	←	←	→	→	→
Current	+	+	+	+	+	+	+
AccelX	0	0	0	-	0	0	0
Bumper							
Action	→	←	←	←	←	←	→
Current	+	+	+	+	+	+	+
AccelX	+	0	0	0	0	-	0
Bumper							
Action	←	←	←	←	←	←	→
Current	++	+	+	+	+	+	+
AccelX	0	0	0	0	0	-	0
Bumper	∗						

successful turning behavior between the walls ($p = 0.02$). The last block is an example for the robot's experience when it collides with the rear wall ($p = 0.001$). After switching to backward movement in time before an imminent collision with the front wall at $t - 4$, it continues to move backwards until it eventually bumps into the rear wall at $t + 1$, triggering the bumper and increasing the current consumption. Since these SMCs result from exploring the properties of the confinement, it is suggested that they correspond to object-related SMCs.

4 Conclusions

The properties of modality- and object-related SMCs suggest that they differ with respect to the temporal context they take into account. Modality-related SMCs reflect the momentary changes in the properties of sensory signals, and give rise to the different qualities of sensory experiences like seeing, hearing, touching etc. Object-related SMCs, in contrast, reflect changes in the sensory information during exploratory processes, and hence, account for sequences of actions and sensory observations that extend over time. They are more specific than modality-related SMCs in the sense that their acquisition and exercise requires a specific configuration of the environment, e.g. the presence of a particular object. Beyond that, we suggest that this type of SMCs actually constitutes what potential objects are for the agent, which then would be defined as repertoires of family-resemblant SMCs.

Our proposed model of SMCs makes the different temporal context sizes explicit. Markov models for the probability distribution of a new action-observation pair conditional on only recent histories of previous actions and observations represent the direct effects that the actions of an artificial agent have on the associated changes of the sensory signal properties. Models in which this probability distribution depends on a longer history of events are more specific for a particular environmental situation, corresponding to object-related SMCs. The robot study shows how both SMC types determine the overt behavior of the agent: While modality-related SMCs let it move in a coordinated and energy-efficient manner, object-related SMCs support its optimal behavior in the specific environmental situation.

It might be interesting to ponder about the implications of extrapolating the idea of a relation between types of SMCs and time scales to contexts beyond object-related SMCs. Behavior is nothing an agent happens to have, but is something that is exercised to bring about specific consequences. These consequences are fed back to the behaving system on a longer time scale. Time-extended actions, like driving to work or preparing a cake for example, are highly structured and involve a similar regularity of sensorimotor correlations like SMCs in the sense of SMCT. It is the precise sequence of turns taken with the car that determines the final destination of the ride, and the specific selection of ingredients and instructions followed from the recipe that define the type of cake made. We suggest, therefore, that the concept of SMCs can be extended to what might be called "intention-related" SMCs that capture these long term regularities in action sequences, and that constitute our conscious experience beyond the time scale of object perception, like "driving to work" or "preparing a cake" in this example. We will develop models of intention-related SMCs and examine the implications of this idea in the future.

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', IST-270212, esmcs.eu.

References

1. O'Regan, J., Noë, A.: A sensorimotor account of vision and visual consciousness. Behavioral and Brain Sciences 24, 939–1031 (2001)
2. Noë, A.: Action in perception. MIT Press (2004)
3. Engel, A.: Directive minds: how dynamics shapes cognition. In: Stewart, J., Gapenne, O., Di Paolo, E.A. (eds.) Enaction: Towards a New Paradigm for Cognitive Science, pp. 219–243. MIT Press, Cambridge (2011)
4. Choe, Y., Yang, H.F., Eng, D.Y.: Autonomous learning of the semantics of internal sensory states based on motor exploration. International Journal of Humanoid Robotics 4, 211–243 (2007)
5. Fine, P., Di Paolo, E.A., Izquierdo, E.: Adapting to Your Body. In: Almeida e Costa, F., Rocha, L.M., Costa, E., Harvey, I., Coutinho, A. (eds.) ECAL 2007. LNCS (LNAI), vol. 4648, pp. 203–212. Springer, Heidelberg (2007)
6. Bovet, S., Pfeifer, R.: Emergence of delayed reward learning from sensorimotor coordination. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 2272–2277 (August 2005)
7. Möller, R., Schenck, W.: Bootstrapping cognition from behavior – a computerized thought experiment. Cognitive Science 32(3), 504–542 (2008)
8. Maye, A., Engel, A.: A discrete computational model of sensorimotor contingencies for object perception and control of behavior. In: 2011 IEEE International Conference on Robotics and Automation (ICRA), pp. 3810–3815. IEEE (May 2011)
9. Pfeifer, R., Scheier, C.: Understanding Intelligence. MIT Press (September 2001)