

Towards an Adaptive Hierarchical Anticipatory Behavior Control System

Oliver Herbort and Martin V. Butz and Joachim Hoffmann

oliver.herbort@stud-mail.uni-wuerzburg.de

{butz,hoffmann}@psychologie.uni-wuerzburg.de

Department of Cognitive Psychology, Röntgenring 11

University of Würzburg

Würzburg, 97070, Germany

Abstract

Despite recent successes in control theoretical programs for limb control, behavior-based cognitive approaches for control are somewhat lacking behind. Insights in psychology and neuroscience suggest that the most important ingredients for a successful developmental approach to control are anticipatory mechanisms and hierarchical structures. Anticipatory mechanisms are beneficial in handling noisy sensors, bridging sensory delays, and directing attention and action processing capacities. Moreover, action selection may be immediate using inverse modeling techniques. Hierarchies enable anticipatory influences on multiple levels of abstraction in time and space. This paper provides an overview over recent insights in anticipatory, hierarchical, cognitive behavioral mechanisms, reviews previous modeling approaches, and introduces a novel model well-suited to study hierarchical anticipatory behavioral control in simulated as well as real robotic control scenarios.

Introduction

The autonomous control of the own body is an essential challenge for any cognitive system. Although established behavioral control in animals and humans seems effortless in every day life, many challenges arise. Due the complex, dynamic, time lagged, noisy, and often nonlinear interactions between body and environment, effective body control in real environments is hard. Movements of different body parts influence each other, clothing change the interactions, muscle forces are state-dependent, etc. Furthermore, sensory information may be unavailable, as for example in darkness, or may be available only after a significant time delay. The mind has to learn these complex, often context-dependent, interactions to be able to induce effective adaptive body control.

The notion that most actions are goal directed and that the goal state is represented before the action is performed is labeled the ideomotor principle, which can be traced back to the 19th century (Herbart 1825; James 1890)). Although behaviorists later questioned this view, it is now widely accepted that behavior is in most cases goal oriented. Hoffmann (1993) emphasized this insight in his theory of anticipatory behavior control, which theorizes that actions are

usually preceded by an anticipatory image of the sensory effects. The image triggers that action(s) that is (are) expected to yield the anticipated effects, considering the current environmental circumstances. Different sensory modalities and sensory aspects can influence action triggering, as, for example, an external effect (such as a tone) or also a proprioceptive effect (such as muscle tension or pressure feedback). To control more complex behavior, actions may be divided into simpler parts. For example, if a piano player wishes to play a tone, the anticipation of the tone causes the anticipation of the feeling of the correct hand position and then the finger pressing the key. Thus, to achieve an overall goal, several successive goals may trigger successive actions.

To be able to generate such complex behavior effectively, hierarchical processes are necessary that generate goals and partition far-reaching goals into suitable sub-goals. However, even if neuroscience shows that brain functions are structured hierarchically (e.g. Poggio & Bizzi 2004), only few computational arguments exist, why such structures are advantageous.

This paper reviews evidence for anticipatory guided human processing and derives design suggestions for cognitive behavior systems. Similarly, we assess evidence for hierarchically structured mechanisms. The gained insights lead us to the development of a simple learning system for studying the potential benefits of hierarchical anticipatory control structures. We introduce the base model and confirm successful behavioral control of a simple arm. In sum, this paper develops anticipatory hierarchically controlled systems that learn effective control structures to guide complex adaptive behavioral patterns.

The remainder of this work is structured as follows. First, we review anticipatory and hierarchical cognitive structures. Next, existing cognitive control models are compared. Finally, we introduce our model revealing its current capabilities, limitations, and potentials. A short discussion concludes the paper.

Anticipatory Hierarchical Structures

In this section, we gather evidence for and benefits of anticipatory and hierarchical structures in learning, behavioral control, and cognition in the broader sense.

Anticipatory Behavior Control

Anticipatory behavior refers to behavior in which currently desired goals precede and trigger the action that usually results in the desired goals. Psychological experiments underline the concept of anticipatory behavior.

A simple experiment confirms the presence of effect representations before action execution. Kunde (2001) paired actions with compatible or incompatible effects, such as the presentation of a bar on the left or on the right compatible or incompatible to a left or right key press. Although the effects were presented only after the key press, reaction times were significantly faster, when the location of the target button and of the visual effect corresponded. Similar effects were found for the modalities of intensity and duration (Kunde, Koch, & Hoffmann 2004). Elsner and Hommel (2001) showed that reaction times also increase if an action is accompanied by a stimulus that does not match with the expected effect, even if this stimulus could be completely ignored to choose the correct response key. In all cases, it is concluded that anticipatory effect representations interfere with an action code or also with an external stimulus. Thus, goal aspects are represented before action execution in terms of at least some of the sensory effects. Interestingly, it has also been shown that humans acquire such action-effect associations much easier than situation-action relations (Stock & Hoffmann 2002).

However, the advantages of such anticipatory behavior remain somewhat obscured. What are the benefits of representing effects before or actually for action execution? Other disciplines provide interesting insights in this respect. Artificial intelligence shows that anticipatory representations enable higher flexibility in learning and decision-making. In reinforcement learning (Sutton & Barto 1998), the DYNA architecture (Sutton 1990) showed that model-based reinforcement learning mechanisms increase flexibility when goals vary or when the environment is partially dynamic. More recent investigations in relational reinforcement learning have shown similar advantages when the flexible propagation of reinforcement learning is required (Kersting, Van Otterlo, & De Raedt 2004).

In control theory, structures capable of predicting future states yield more powerful controllers. For example, a forward model that predicts the consequences of actions may be used to correct errors in advance (Miall *et al.* 1993). The concept of combining sensory and predictive information to compensate for unavailable, delayed, or highly noisy sensory feedback is made most explicit in the widely applied Kalman filter (Haykin 2002; Kalman 1960). Neuroscientific studies indicate that Kalman filtering-like structures exist in the cerebellum (Barlow 2002). Additionally, it was shown that inverse models (IMs), that directly determine the action necessary to obtain a desired goal, result in efficient adaptive controllers (Kawato, Furukawa, & Suzuki 1987).

Thus, cognitive psychology and neuroscience suggest that anticipations are important for effective adaptive learning systems. Artificial intelligence and control theory have shown that anticipatory structures improve learning speed and reliability, behavioral flexibility and control, and sensory robustness, resulting in effective goal-directed systems.

Hierarchies for Learning and Control

Besides the anticipatory indicators, studies and models suggest that cognitive information is processed hierarchically. Powers (1973) already stressed the importance of hierarchies in behavioral control and consequent computational models of cognitive systems. Just recently, Poggio and Bizzi (2004) pointed out that hierarchical structures are very likely the key not only to sensory processing but also to motor control. Available hierarchical models in vision (Giese & Poggio 2003; Riesenhuber & Poggio 1999) are suggested to be extended to motor control. Hierarchical top-down influences showed to have advantageous structuring effects (Rao & Ballard 1999).

Computational motor control models showed advantages of hierarchical structures. Considering the hierarchy of the musculoskeletal system, the spinal cord, and a controller in the CNS at the top, Loeb, Brown and Cheng (1999) demonstrated that the spinal cord is able to counter most perturbations on its own. However, the spinal cord also receives task-depending input from the CNS to adjust its behavior. Thus, the spinal cord makes the control task easier for the CNS because not every single muscle has to be addressed. It is sufficient to set an overall strategy to deal with most perturbations.

Hierarchical processing models were proposed by Kawato, Furukawa and Suzuki (1987), who applied a hierarchical controller to a robot arm. The lowest level contains a simple PD-controller that can in principle handle any task. The controller is not very efficient, because the delayed feedback results in a slow control process. A second layer improves performance. As soon as a direct model of the plant is learned, it updates the control signal using the expected feedback, which is available much faster than the actual, sensory feedback. However, it is still necessary to adjust the signal iteratively. A third level consists of an *inverse model* (IM) that calculates a control signal for any given goal. When the IM is accurate, the controller selects a feasible control signal instantly. In case of a failure, the lower levels induce the (slower and less effective) control. The more accurate the models in the higher levels, the more they influence the control signals.

Despite the ubiquitous hints on the importance of hierarchical processing and the first model from Kawato and colleagues, it remains somewhat unclear why hierarchies are advantageous. One advantage may be the rather general decomposability of our environment due to time and space constraints (Gibson 1979; Simon 1969). Computational advantages can be found in artificial intelligence studies.

Reconsidering reinforcement learning, it has become clear that hierarchical processing mechanisms are mandatory for effective reward propagation and flexible learning (Barto & Mahadevan 2003). Hierarchical structures are formed that can trigger options, that is, abstract action representations including goals. Most recent publications have shown that such hierarchical representations may be learned by using simple statistics of the environment searching for significantly independent sub-structures (Butz, Swarup, & Goldberg 2004; Simsek & Barto 2004).

Thus, hierarchical control enables the discovery and re-

presentation of more distant and abstract dependencies as well as increases flexibility in behavioral learning and decision making, as well as in sensory processing at different levels of abstraction in time and space.

Merging Both

As we have seen, cognitive processing is guided by anticipations that improve sensory processing and behavioral control. Hierarchies yield more flexible representations for anticipatory learning and behavior. The observations suggest that the combination of anticipatory and hierarchical structures may be a promising approach to understand and model human motor learning and control.

The review also suggests several requirements for a cognitive controller. First, the controller must represent a goal in terms of desired sensory inputs. Partial, underspecified, and even contradicting goals may be represented in different sensory modalities. Second, goal representations should not only be modular but also hierarchical. Higher level goal representations are usually more abstract in time and space and trigger lower level, more concrete, sensory dependent goal representations. Third, the representations should be learned by interacting with the environment. Learning architecture and learning biases, however, are provided in advance.

Biological Plausibility

To model motor learning and control, not only functional constraints have to be taken into account. Additionally, the structure of the motor control system and the mechanisms that modify the structure during learning should be biologically plausible.

Neural networks are considered realistic models of knowledge representations in the brain. In the case of multi-layer neural networks, this plausibility holds only for an already learned network but not for the training mechanisms. For single layer networks, the Hebbian learning rule (Hebb 1949) provides a biologically plausible learning algorithm. It states that connections between neurons are strengthened, if both neurons are excited at the same time and weakened otherwise. Thus, it forms a basis for associative learning. Unfortunately, Hebbian learning only works for single layer neural networks that can only compute linear separable functions that are too simple for motor control. However, to overcome this problem actions and goals can be represented in a clustered form that divides the learning space in small subspaces. The division enables the storage of non-linear goal-action mappings.

Before we introduce our model, which can satisfy these constraints, we review other related systems.

Cognitive Movement Controllers

Numerous computational models for motor learning and control have been proposed. Most of them address specific stages of movement generation, for example trajectory formation (Cruse, Steinkühler, & Burkamp 1998; Hirayama, Kawato, & Jordan 1993) or coordinate transformation (Salinas & Abbott 1995). Others are tracking reference signals, relying on IMs and feedback controllers (Kalveram 2004;

Kalveram *et al.* 2005; Kawato, Furukawa, & Suzuki 1987), which might be combined in a single control structure (Stroeve 1996; 1997). Some approaches gate a number of single control structures to be able to quickly adapt to changing limb properties (Haruno, Wolpert, & Kawato 2001; Wolpert & Kawato 1998) or to combine motor primitives (Berthier *et al.* 1992).

While each model has interesting properties on its own, none match all the suggested cognitive system requirements. The described hierarchical model of Kawato, Furukawa and Suzuki (1987) contains three different levels but does not accept goals in arbitrary modalities. Other controllers (Cruse, Steinkühler, & Burkamp 1998) accept underspecified goals but do not include hierarchical layers. Many models contain neural networks that learn by cognitively implausible mechanisms, such as back-propagation. Our model intends to bridge the respective drawbacks effectively creating a hierarchical, anticipatory cognitive model that is suitable to process any goal representation flexibly and hierarchically.

A Hierarchically Anticipatory Model of Motor Control

We devise a new computational model for motor learning and control. The central part of the model is the controller that can transform any goal, represented in any available sensor modalities, into action signals that move the body to or at least towards a position, in which the desired sensory effects are (usually) perceived. This goal action mapping has to be learned by the controller while interacting with the environment.

Inverse Models and Sensory Representation

A structure that transforms goals into actions usually needs to represent a complex non-linear function. To do this with a single layer neural network, the learning space can be divided into small parts using Radial Basis Functions (RBF). A sensory signal then is not represented by a single neuron with an activation that correlates with a variable of the body configuration, but it is represented by an array of single neurons that represent a specific range of the possible values of a variable. For example, a joint angle is not encoded by a neuron that has a growing firing rate with growing limb extension but by many different neurons. Each of these neurons is only activated, if the joint angle is in a specific range. The range of values for which a neuron is activated is called its *receptive field*. This kind of representation fits well to electro-physiological data obtained from measuring the correspondence of single cell activity in the motor cortex and movement patterns (Georgopoulos 1995).

Actions are represented in the same fashion, when observed by the controller. However, this representation is not very likely to exist in the periphery (motoneurons, proprioceptions, etc). Thus, two transformations have to be done. To encode a perception like a joint angle into an array of neurons, the activation of every neuron has to be determined. The receptive field of a neuron is characterized by its center and width. The center is the variable value that will maximally excite the neuron. The width is given by the range

of values that will cause any activation. In the model, the activation of a neuron is calculated by applying a Gaussian distribution function to the distance between the center of the receptive field and the measured value (using half the distance to the next center of a receptive field as standard deviation).

To reverse these transformations, a winner-takes-all mechanism converts the activity of many neurons into a single signal. Thereby, the activation of the output signal is set to the center of the receptive field of the neuron with the highest activation.¹

To learn the neural network, actions have to be associated to their effects, according to the situation. This is done by strengthening the connections between neurons that encode a situation and the effect of an action and the neurons that encode actions, if both are activated at the same time. After learning, the network can choose an action that will produce the desired sensory effect in a specific situation. Thus, an inverse model can be learned in a biologically plausible way by interacting with the environment.

Generalization

The neural network presented above raises several questions. At a first glance, the controller seems incapable of generalization, that is, performing actions or reaching goals that had not been encountered, yet. This shortcoming is partially solved by the representation of the information. If the receptive fields of the neurons are wide, they are even activated, if the sensory signal is not in the close vicinity of the center of the receptive field. Hence, the network will contain information about what to do also if the currently desired sensory state has not been reached during learning. This spatial generalization capability does not interfere with knowledge over experienced movements, because activations due to generalization are comparatively low.

Naturally, it takes some time until an action has a noticeable effect in the environment. To account for this, the neural network relates actions to the sensory state that is perceived a few moments later. However, this restricts the inverse model to store action-effect relations that are encountered in a short time interval.

Consider the movement of an arm that needs 300ms to move from a relaxed to a fully extended position. If the neural network encodes only the effects that an action has after 100ms, it will not contain information about how to fully extend the relaxed arm, because it cannot be done in 100ms. This problem can be reduced by introducing a second kind of generalization (temporal generalization) that not only relates more distant points in sensory space to certain actions, but also sensory effects that occur after a longer time interval. Again, to reduce interference with actually observed movements, the connections between sensory effects that are produced later in time are exponentially weaker.

¹To enable more fine grained output signals, the neighbors of the winning neuron are also taken into account, weighing their influence according to their activation level. If there are several winning neurons, one winner is chosen randomly.

These two types of generalization allow the network to control an arm effectively and reliably.

Model Evaluation

To test the feasibility of our approach, we evaluate the performance of a single IM on a simple 1-dof-arm model. The arm is dampened and a restoring force pulls it to its initial position. To move the arm, a motor signal is proportionally transformed into a torque that is applied to the joint. Thus, applying a constant torque signal moves the arm to a certain equilibrium position after some oscillations.

The IM is capable of applying a torque to produce a specific desired effect that comprises the joint angle and its velocity. Receptive fields of sensory states (desired or actual) are distant $0.1rad$ ($1rad = 5.7deg$) for joint angle and on average $1\frac{rad}{s}$ for velocities, with a higher resolution for small velocities. For the encoding of the action (torque) receptive fields are distant $10\frac{rad}{s^2}$. A change of torque of about 1 moves the arm to a new equilibrium position about $0.067rad$ distant from the current one.

The network is trained by applying a new random torque between $-150\frac{rad}{s^2}$ and $150\frac{rad}{s^2}$ for 50ms, 100ms, 150ms or 200ms. This causes the arm to move randomly in the range of $-1rad$ to $1rad$. To test the controller, the arm has to move from 5 different starting angles ($-0.66rad, -0.33rad, 0.0rad, .33rad, .66rad$) to 21 different targets ($-1.0rad, -0.9rad, \dots, 1rad$). In all cases, the desired velocity is set to 0. During a reaching movement, the controller sets a new torque every 50ms. A movement is considered finished, if the joint angle does not move more than $0.02rad$ within 250ms.

To test whether the IM benefits from generalization capabilities, both spatial and temporal generalization were varied in a 2×2 -design with 20 independently trained IMs in each of the four groups. Spatial generalization was manipulated by altering the width of the receptive fields of the neurons. The high spatial generalization condition had receptive fields that measured 3 times the distance to the next center of a receptive field in diameter. The diameter was reduced by a third in the low spatial generalization condition. Temporal generalization was manipulated by allowing the controller to relate actions to effects that occurred up to 100ms (low) or up to 150ms (high) after the action had been performed. In

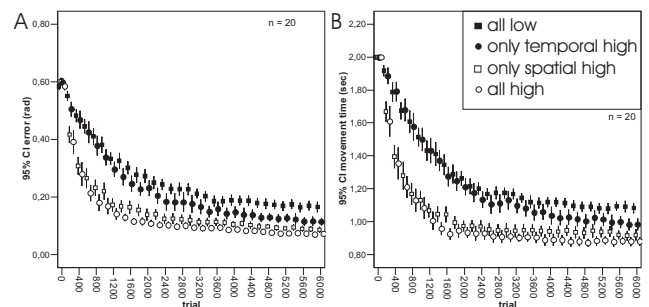


Figure 1: Spatial and temporal generalization yield advantages for accuracy (A) and movement speed (B).

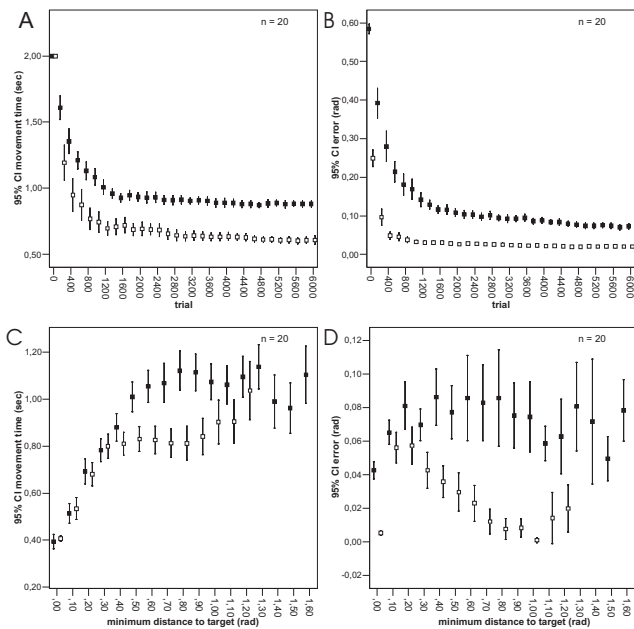


Figure 2: A range of acceptable end positions (white squares) can be reached faster (A) and with less error (B) than an exactly specified goal angle (black squares). This also holds, if movement time (C) and accuracy (D) of movements to goal ranges are compared to movement that try to reach the nearest point of a given goal range.

both cases, effects that occurred later had an activation level of about 1% of the initial activation level. Figure 1 shows that both spatial and temporal generalization capabilities allow increased accuracy (A) and faster movements (B). Split-plot ANOVAs revealed significant main effects between the groups for movement time and accuracy (both $p < .01$). The interactions between the groups did not reach significance.

A feature of the RBF-like representation is, that not only precise goal coordinates can be desired, but also ranges of acceptable goal positions. It is not always necessary to reach an exact end position. To test if it is advantageous to give a wider goal, reaching to an exact position was compared to reaching to position anywhere within a range of 0.8rad . Each group consists of 20 IMs that learned independently. Figures 2 A and B show, that movements to exact positions (black squares) are slower (A) and less exact (B) than movements to wide goal ranges (white squares). Note, that the error in the goal-range-condition is calculated as the distance to the nearest joint angle within the range. Split-plot ANOVAs confirmed both results ($p < .01$). This also holds, when the movement to a goal range is compared to a movement to the exact point within the range that is closest to the initial position. On average, movements to any point in the goal range are faster (C) and yield less errors (D) than movements to the nearest point in the goal range. ANCOVAs that controlled for the distance to the nearest possible goal confirmed both conjectures ($p < .01$).

Towards an Hierarchical Controller

The data presented above shows that it is possible to train a single layer neural network in a biologically plausible manner to control a simple arm model. As outlined above, generalization mechanisms can be used to make the IM reach targets that have not been reached before or would be out of scope because the distance between the initial sensory state and the desired sensory state is to large.

We claimed, that a controller should be able to process many different modalities. A single IM that relates all kinds of sensory inputs (from proprioceptions to distal effects) would require a huge neural network structure and challenge its temporal generalization capabilities, because it would have to relate complex muscle activation patterns with events that happen not very often and that can be produced in many different ways. Additionally, the network would have to learn any actions by rote learning.

Consider switching on a lamp. To change the sensory input that encodes brightness by pushing the switch, a long sequence of muscle activations has to be carried out. To learn this sequence by random movements may take a very long time. The problem would be easier to solve, if many

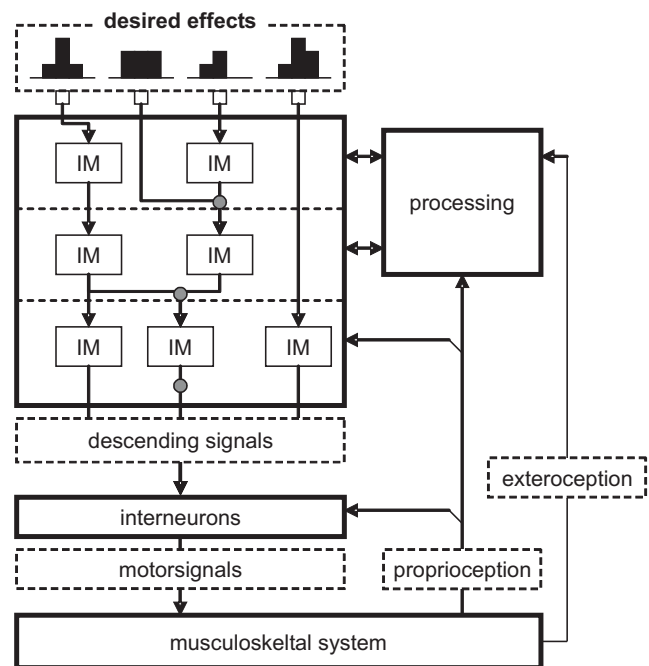


Figure 3: The left side of the drawing shows how the desired effects are transformed into motor signals using a hierarchy of inverse models (IM). The motor signals cause changes in the arm. The changes are fed back to the controller (right side). Perceptions can be used directly or after processing.

different IMs were involved. An IM that relates brightness to hand position might well learn that the light goes on if the hand reaches the position where the switch is. A second IM that stores which muscle activation patterns are needed to reach a specific hand position could then be used

to move the hand there. Additionally, other tasks that need the hand to reach a specific position can be easily learned because they can share the knowledge of the handposition-to-muscle-activation-IM. Thus, a complex goal like switching on the light can be transformed into a more concrete desired effect like a hand position that can be more easily transformed into actual muscle activations.

The scheme of such a structure is outlined in figure 3. More complex or abstract goals like producing a tone with a piano or switching on a lamp are converted by IMs into subgoals of a different modality. The subgoals are then considered as desired effect by IMs on a lower layer of the hierarchy and are thus transformed to sub-subgoals and so on. It should be possible to not only give goals on the highest level, but also to bypass layers and desired more concrete goals like a specific hand position. The IMs lowest layer then calculates signals that are sent to lower motor centers or the spinal cord, where they are transformed into muscle activations after some processing. The muscle activation may cause the body to move and thus evoke new proprioceptions and exteroceptions that are available to the controller directly (like muscle tension measured by the Golgi tendon organs) or after computing some abstract representations such as the hand coordinates given visual information.

Thus, a hierarchy of IMs can be used to enhance learning and control. First, the learning space is divided into sub spaces, so that it is not necessary to associate concrete actions to any sensory effects but it is also possible to connect new sensory effects (like brightness) to already known sensory effects (like hand position). Second, by storing the information in different IMs, each IM can be accessed directly and thus any goal modality can easily be used as input.

Model Capabilities, Potentials, and Challenges

The data showed, that the single IMs can learn to control a single arm by mere associative learning. Thereby, the controller uses efficient activation sequences to reach a target.

A major point of critique on the current model may be the method used for learning the IM. Since the method is not goal-directed, the mapping is nowhere guaranteed to converge to the optimum (Jordan & Rumelhart 1992). However, we believe that it is not necessary to obtain an optimally accurate mapping between action and effects in the general sense. Action execution is usually noisy and easily perturbed so that sensory feedback control is expected to be generally necessary to reach a precise goal. The IMs presented can also be used for closed loop control (Herbort, Butz, & Hoffmann 2005).

Another concern is that the chosen RBF encoding is not very suitable for generalization. Using different layers of RBFs with a combination of larger and smaller receptive fields may solve this problem. However, the encoding also has advantages. The representation facilitates dealing with uncertainty (Knill & Pouget 2004) and allows very flexible goal representations. A goal does not need to be exactly specified but a range of acceptable goal states or goal features can be presented to the network. This capability increases flexibility, which is advantageous for control (Todorov & Jordan 2002). Additionally, the representation

allows the encoding of many-to-many relationships. A final RBF-related concern may be the curse of dimensionality and the consequently exploding number of RBF neurons. However, techniques exist that can reduce the number of neurons by adapting RBF sizes to the demands of the action-effect function (see (Butz 2005) for one potential mechanism). Additionally, separating comparatively independent parts of the sensory space in different networks can reduce the amount of required neurons (Urban, Bruessler, & Gresser 1998).

A big challenge arises considering the need to learn and execute motor programs. Currently the system state only changes, if the desired effects or sensory inputs change. Thus, very fast or complex movements are not possible. Two ways exist to integrate motor programs. First, it has been shown that neural circuits exist in the spinal cord of animals that generate specific motor signals to coordinate simple rhythmic behavior, like walking or swimming (Dietz 2003). Thus, the model of the spinal cord could be extended to include such rhythmic pattern generators. Additional representations would be necessary to code the behavior caused by the pattern generators, such as representations of walking or moving forward, to be able to address the behavior with anticipations. A second way to include motor programs would be to delegate this task to higher structures that send continuously changing, desired effects to the controller. The combination of both features may be able to learn rhythmic behavior combined with consecutive behavioral pattern changes, as appropriate.

In this paper we only presented results for one single controller. Experiments are in progress combining multiple controllers as outlined above. Two approaches need to be distinguished: parallel, modular combinations and hierarchical, abstracting combinations. Shadmehr and Brashers-Krug (1997) showed that human subjects are able to store many different controllers for different situations. For example, one controller could be trained for moving light objects and another for heavy objects. The weighted combination of both controllers then enables fast adaptation to specific situations. This feature can be added by using an array of controllers that are experts for a specific situation and are weighted accordingly (Haruno, Wolpert, & Kawato 2001). Hierarchically connected IMs might prove advantageous when different objects need to be moved. Although different weighting of lower level IMs is necessary to calculate descending commands from the desired joint angles, the relationship between external coordinates and joint angles stay constant. Thus, only parts need to be adapted to the current situation. Additionally, longer time delays in higher layers may be compensated for by lower level control structures. On the other hand, the different times integrated by different models may be used to facilitate more complex, longer term movements.

Besides the combination and extension of IMs, strongly noisy signals will require more elaborate processes. Forward models can be included in the processing of the sensory inputs to bridge temporary misperceptions, sensory failure, or noisy sensory inputs akin to Kalman filtering.

Summary and Conclusion

This paper has reviewed indications and benefits of anticipatory mechanisms and hierarchical structures in control processes. Both mechanisms are involved in human motor learning and control. While anticipatory mechanisms lead to direct action selections in inverse models and effective filtering mechanisms in forward models, the modular and hierarchical combination of such models promises to yield a more effective environmental representation increasing behavioral flexibility, adaptivity, and decision making.

The gathered potentials of combining both mechanisms into artificial cognitive systems promise fruitful future research. The proposed model provides a novel, integrative approach for studying such combinations. The generality of the proposed associative structures enables direct modular and hierarchical combinations. Future research will investigate the suitability and extendibility of our approach for the simulation of efficient cognitive learning systems in simulated and real robotic environments. Moreover, future research will further study the benefits of hierarchical, anticipatory behavior control for learning, behavior, and cognition in general, using and extending the proposed model.

Acknowledgments

This work was sponsored by the German research foundation (DFG) under grant DFG HO1301/4-3 as well as by the European commission contract no. FP6-511931.

References

- Barlow, J. S. 2002. *The cerebellum and adaptive control*. New York, NY: Cambridge University Press.
- Barto, A. G., and Mahadevan, S. 2003. Recent advances in hierarchical reinforcement learning. *Discrete Event Dynamic Systems* 13:341–379.
- Berthier, N. E.; Singh, S. P.; Barto, A. G.; and Houk, J. C. 1992. A cortico-cerebellar model that learns to generate distributed motor commands to control a kinematic arm. *Advances in Neural Information Processing Systems* 4:611–618.
- Butz, M. V.; Swarup, S.; and Goldberg, D. E. 2004. Effective online detection of task-independent landmarks. IliGAL report 2004002, Illinois Genetic Algorithms Laboratory, University of Illinois at Urbana-Champaign.
- Butz, M. V. 2005. Kernel-based, ellipsoidal conditions in the real-valued XCS classifier system. *GECCO 2005: Genetic and Evolutionary Computation Conference: Volume 2* 1835–1842.
- Cruse, H.; Steinkühler, U.; and Burkamp, C. 1998. MMC - a recurrent neural network which can be used as manipulable body model. In Pfeifer, R.; Blumberg, B.; Meyer, J.-A.; and Wilson, S., eds., *From Animal to Animats 5*, 381–389. Cambridge, MA: MIT Press.
- Dietz, V. 2003. Spinal cord pattern generators for locomotion. *Clinical Neurophysiology* 114:1379–1389.
- Elsner, B., and Hommel, B. 2001. Effect anticipations and action control. *Journal of Experimental Psychology* 27(1):229–240.
- Georgopoulos, A. P. 1995. Current issues in directional motor control. *Trends Neuroscience* 18(11):506–510.
- Gibson, J. J. 1979. *The Ecological Approach to Visual Perception*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Giese, M. A., and Poggio, T. 2003. Neural mechanisms for the recognition of biological movements. *Nature Reviews Neuroscience* 4:179–192.
- Haruno, M.; Wolpert, D. M.; and Kawato, M. 2001. Mosaic model for sensorimotor learning and control. *Neural Computation* 13(10):2201–2220.
- Haykin, S. 2002. *Adaptive filter theory*. Upper Saddle River, NJ: Prentice Hall, 4th edition.
- Hebb, D. O. 1949. *The Organization of Behavior*. New York: John Wiley.
- Herbart, J. F. 1825. *Psychologie als Wissenschaft neu gegründet auf Erfahrung, Metaphysik und Mathematik. Zweiter analytischer Teil*. August Wilhelm Unzer.
- Herbert, O.; Butz, M. V.; and Hoffmann, J. 2005. Towards the advantages of hierarchical anticipatory behavioral control. In Opwis, K., and Penner, I.-K., eds., *Proceedings of the KogWis05. The German Cognitive Science Conference*, 77–82. Basel: Schwabe.
- Hirayama, M.; Kawato, M.; and Jordan, M. I. 1993. The cascade neural network model and a speed-accuracy trade-off of arm movement. *Journal of Motor Behaviour* 25(3):162–174.
- Hoffmann, J. 1993. *Vorhersage und Erkenntnis*. Göttingen, Germany: Hogrefe.
- James, W. 1890. *The principles of psychology*, volume 1. New York: Holt.
- Jordan, M. I., and Rumelhart, D. E. 1992. Forward models: Supervised learning with a distal teacher. *Cognitive Science* 16:307–354.
- Kalman, R. E. 1960. A new approach to linear filtering and prediction problems. *Transactions of the ASME-Journal of Basic Engineering* 82(Series D):35–45.
- Kalveram, K. T.; Schinauer, T.; Beirle, S.; Richter, S.; and Jansen Osmann, P. 2005. Threading neural feedforward into a mechanical spring - how biology exploits physics in limb control. *Biological Cybernetics*.
- Kalveram, K. T. 2004. The inverse problem in cognitive, perceptual and proprioceptive control of sensorimotor behaviour: towards a biologically plausible model of the control of aiming movements. *International Journal of Sport and Exercise* 2:255–273.
- Kawato, M.; Furukawa, K.; and Suzuki, R. 1987. A hierarchical neural-network model for control and learning of voluntary movement. *Biological Cybernetics* 57:169–185.
- Kersting, K.; Van Otterlo, M.; and De Raedt, L. 2004. Bellman goes relational. *Proceedings of the Twenty-First International Conference on Machine Learning (ICML-2004)* 465–472.
- Knill, D. C., and Pouget, A. 2004. The bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences* 27(12):712–719.

- Kunde, W.; Koch, I.; and Hoffmann, J. 2004. Anticipated action effects affect the selection, initiation, and execution of actions. *The Quarterly Journal of Experimental Psychology. Section A: Human Experimental Psychology* 57:87–106.
- Kunde, W. 2001. Response-effect compatibility in manual choice reaction tasks. *Journal of Experimental Psychology: Human Perception and Performance* 27(2):387–394.
- Loeb, G. E.; Brown, I.; and Cheng, E. 1999. A hierarchical foundation for models of sensorimotor control. *Experimental Brain Research* 126:1–18.
- Miall, R. C.; Weir, D. J.; Wolpert, D. M.; and Stein, J. F. 1993. Is the cerebellum a smith predictor? *Journal of Motor Behavior* 25(3):203–216.
- Poggio, T., and Bizzi, E. 2004. Generalization in vision and motor control. *Nature* 431:768–774.
- Powers, W. T. 1973. *Behavior: The Control of Perception*. New York: Aldine de Gruyter.
- Rao, R. P. N., and Ballard, D. H. 1999. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience* 2(1):79–87.
- Riesenhuber, M., and Poggio, T. 1999. Hierarchical models of object recognition in cortex. *Nature Neuroscience* 2(11):1019–1025.
- Salinas, E., and Abbott, L. F. 1995. Transfer of coded information from sensory to motor networks. *Journal of Neuroscience* 15(10):6461–6474.
- Shadmehr, R., and Brashers-Krug, T. 1997. Functional stages in the formation of human long term motor memory. *Journal of Neuroscience* 17(1):409–419.
- Simon, H. A. 1969. *Sciences of the Artificial*. Cambridge, MA: MIT Press.
- Simsek, O., and Barto, A. G. 2004. Using relative novelty to identify useful temporal abstractions in reinforcement learning. *Proceedings of the Twenty-First International Conference on Machine Learning (ICML-2004)* 751–758.
- Stock, A., and Hoffmann, J. 2002. Intentional fixation of behavioral learning or how R-E learning blocks S-R learning. *European Journal of Cognitive Psychology* 14(1):127–153.
- Stroeve, S. 1996. Linear combined feedback and feed-forward control of a musculoskeletal system. *Biological Cybernetics* 75:73–83.
- Stroeve, S. 1997. A learning feedback and feedforward neuromuscular control model for two degrees of freedom human arm movements. *Human Movement Science* 16:621–651.
- Sutton, R. S., and Barto, A. G. 1998. *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Sutton, R. S. 1990. Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. *Proceedings of the Seventh International Conference on Machine Learning* 216–224.
- Todorov, E., and Jordan, M. I. 2002. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience* 5(11):1226–1235.
- Urban, J.-P.; Bruessler, J. L.; and Gresser, J. 1998. Neural networks for visual servoing in robotics. In Jain, L. C., and Fukuda, T., eds., *Soft Computing for Intelligent Robotic Systems*. Springer. 81–111.
- Wolpert, D. M., and Kawato, M. 1998. Multiple paired forward and inverse models for motor control. *Neural Networks* 11(7-8):1317–1329.