Learning and Adaptation of Sensorimotor Contingencies: Prism-Adaptation, a Case Study

Gert Kootstra[1], Niklas Wilming[2], Nico Schmidt[3], Mikael Djurfeldt[4], Danica Kragic[1] and Peter König[5]

[1] CAS-CVAP, CSC, Royal Institute of Technology (KTH). {kootstra,dani}@kth.se
[2] Institute of Cognitive Science, University of Osnabrück. nwilming@uos.de
[3] AI Lab, University of Zürich. nico.schmidt@uzh.ch
[4] PDC, Royal Institute of Technology (KTH). mdj@kth.se
[5] Institute of Cognitive Science, University of Osnabrück; Department of Neurophysiology and Pathophysiology, University Medical Center Hamburg-Eppendorf. pkoenig@uos.de

Abstract. This paper focuses on learning and adaptation of sensorimotor contingencies. As a specific case, we investigate the application of prism glasses, which change visual-motor contingencies. After an initial disruption of sensorimotor coordination, humans quickly adapt. However, scope and generalization of that adaptation is highly dependent on the type of feedback and exhibits markedly different degrees of generalization. We apply a model with a specific interaction of forward and inverse models to a robotic setup and subject it to the identical experiments that have been used on previous human psychophysical studies. Our model demonstrates both locally specific adaptation and global generalization in accordance with the psychophysical experiments. These results emphasize the role of the motor system for sensory processes and open an avenue to improve on sensorimotor processing.

Keywords: Sensorimotor contingencies, prism-adaptation, motor learning/adaptation, body maps, inverse kinematics

1 Introduction

Humans adapt easily to changes in sensorimotor coordination during development and adulthood. A remarkable demonstration is adaptation to prism glasses that displace the visual field horizontally by a constant angle. Despite such drastic changes of a sensorimotor relationship, eye-hand coordination quickly adapts [6].

This is, however, not a passive process but requires active exploration [4, 5]. Prism adaptation is specific to the involved body parts [7, 15] and actions [1, 7, 16]. The adaptation to changed sensorimotor dependencies therefore appears to crucially depend on the ability to relate one's own motor actions to their observed sensory consequences. This bears a striking resemblance to the concept of sensorimotor contingencies (SMC) [11], which also stresses the importance of action for perception [3]. The case of prism adaptation is well investigated and therefore may serve as a prime example for investigating how the brain achieves "mastery of a sensorimotor contingency" [11].

When prisms are donned, subject's pointing movements are offset due to the visual displacement. However, with repeated movements this offset diminishes and original

performance restores. When subsequently the prisms are removed, pointing movements are offset in the opposite direction. This *aftereffect*, i.e. the difference in pointing between pre- and post-exposure, is a convenient measure of the degree of adaptation.

The adaptation is thought to combine two separate processes: Recalibration and realignment [15]. Recalibration is believed to utilize a cognitive learning strategy that quickly reduces pointing errors. The effects of recalibration are local to specific actions. Recalibration is quantified by the aftereffect measured with the identical movement that was carried out during the prism exposure phase [15]. Realignment, in contrast, is thought to be an automatic process that aligns, for example, visual and proprioceptive maps. It reveals a global generalization of adaptation to new pointing targets and actions. Realignment is measured by actions not practiced during prism exposure.

Redding and Wallace observe that the realignment effect is modulated by different kinds of feedback during exposure. When participants could see their own hand-movement (concurrent feedback) during exposure, Redding and Wallace [15] observed a small shift of "visual straight ahead" and a large shift of "proprioceptive straight ahead". This pattern was reversed when participants could only see the end-position of their hand (terminal feedback) during exposure. Related to this, Redding and Wallace found that the aftereffect generalized differently to targets not shown during exposure. Specifically, concurrent feedback produced aftereffects that increased for targets in the direction of the prismatic shift, whereas they decreased for terminal feedback. Both results are explained by two different references: During concurrent feedback the visual system acts as a reference and the proprioceptive system is aligned to it and during terminal feedback the situation is reversed.

In this paper, we investigate properties of recalibration and realignment and their dependence on different types of feedback. For this purpose, a computational model of prism adaptation is developed for the control of a simulated robotic arm and subjected to the identical experiments focused on eye-hand coordination of the previous psychophysical study [15]. Specifically, we test the hypothesis that the different forms of adaptation aftereffects can be described in a unified approach based on the concept of sensorimotor contingencies. We take as our starting point that adaptive agents have to relate changes in perception to their own actions. At the same time we strive for a computational description of recalibration and realignment to foster our understanding of sensorimotor adaption and for facilitating the development of versatile robotic agents. The ability to quickly learn eye-hand coordination from own experience makes manual calibrations redundant, which is interesting, for instance, for robotic grasping [12].

2 Computational model of prism adaptation

The contingency between sensory and motor signals involved in pointing movements is captured by forward and inverse-kinematic models. The forward model defines the position of the effector, e.g. the hand, based on the joint angles of the arm. The inverse-kinematic model provides the joint configuration of the arm necessary to reach a specific target, and thereby allows direct control of a robot. Motivated by [2, 10] and in accordance with the SMC theory [11] we use a differential inverse-kinematic model that generates the action necessary to reach a desired *change* in visual position of the

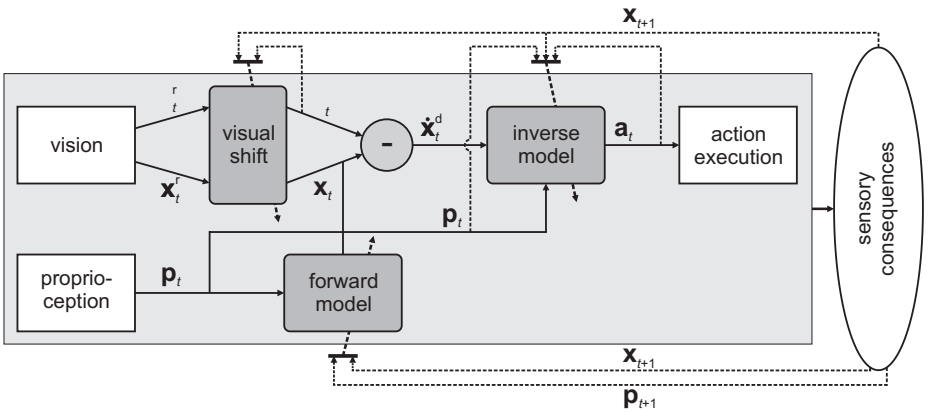

Fig. 1. Theoretical model, including three sub-modules: the visual shift, the forward model, and the inverse model. Based on visual and proprioceptive observations, the model estimates the action to get the end effector to the target pose. All three sub-modules learn on-line using the sensory consequences of the executed action as observed by the system itself. Information flow for control is indicated by solid arrows, and the flow for learning by dashed arrows. The variables are explained in the text.

end-effector. Given a current joint configuration, $\mathbf{q}$, and a desired *change* in pose of the end effector, $\dot{\mathbf{x}}$, the necessary *change* in joint angles, $\dot{\mathbf{q}}$, is determined by the model, expressed by the mapping $(\mathbf{q}, \dot{\mathbf{x}}) \mapsto \dot{\mathbf{q}}$. For small enough changes, the optimal gradient $\dot{\mathbf{q}}$ for any given state $\mathbf{q}$ is well defined and resolves ambiguities inherent in inverse-kinematics.

When the robotic system is subject to changes, for instance, due to mechanical wear or damage or adaptation to prisms, the forward and inverse mappings are not fixed, but changes over time. This calls for adaptive systems that can learn and adapt the inverse-kinematics model from own experience [9]. Others have used, for instance, locally-weighted projection regression [2, 17], Gaussian-process regression (GPR) [13], and local Gaussian-process regression [10]. Here we use Gaussian-process regression (GPR) to learn the inverse kinematics, since it has the best reported performance [13, 10] at the cost of higher computational load. But note, that the GPR can be easily substituted by other regression methods, if real-time performance is more of an issue.

In total, our model consists of three sub-modules: the forward model, the inverse model, and the visual-shift model (see Figure 1). In the following, we give the model details of all components in turn.

2.1 Forward model and inverse model

Based on the proprioceptive information $\mathbf{p}_t$, the *forward model* gives an estimation of the pose of the end effector in visual coordinates $\mathbf{x}_t$ at time t (see Figure 1). This estimation is used when the end effector cannot be observed visually. The *inverse model* provides an estimate of the necessary action, $\mathbf{a}_t$, for instance a change in joint angles,

based on $\mathbf{p}_t$ and the desired change in pose, $\dot{\mathbf{x}}_t^{\mathbf{d}}$, which is determined based on the difference between the target and the end effector, $\dot{\mathbf{x}}_t^{\mathbf{d}} = \tau_t - \mathbf{x}_t$ (see Figure 1).

Using GPR, an estimation is made of the functions $\mathbf{a}_t = f(\mathbf{p}_t, \dot{\mathbf{x}}_t^{\mathbf{d}})$ for the inverse and $\mathbf{x}_t = g(\mathbf{p}_t)$ for the forward model, based on a set of training examples, which the system acquires from own experience. GPR has a few hyperparameters, including $\{\lambda_1^2, \ldots, \lambda_D^2\}$, which are the characteristic length-scales of the different dimensions of the squared-exponential kernel, where D is the dimensionality of the input to the GPR. The hyperparameters are learned from data by maximizing the marginal likelihood. For details on GPR, we refer to [13].

Recency effect A problem with a standard GPR implementation of the forward and inverse models is that adaptation to changed SMCs is slow. When prism glasses are donned, the old training samples contribute as strong to the estimation as the new samples, which results in a rather slow adaptation. To increase speed of adaptation and to intensify the aftereffect, we include a forgetting mechanism using a recency effect, such that more recent training data make a stronger contribution to the estimation. Specifically, we add a time dimension to the input, resulting in input $\mathbf{z}^{\mathrm{i}} = \{\mathbf{p}, \dot{\mathbf{x}}, t\}$ for the inverse, and $\mathbf{z}^{\mathrm{f}} = \{\mathbf{p}, t\}$ for the forward model, both with a time constant λ_T. This characteristic length-scale parameter is not included in the optimizing of the hyperparameters, but instead used as a free parameter to control the recency effect, which is strong for low values of λ_T, whereas for $\lambda_T \to \infty$, the effect is absent.

Execution and on-line learning The Gaussian process regressors that implement the inverse and forward models are continuously updated while the robot performs its task, thus on-line learning the eye-hand coordination. In execution, the GPRs are used to predict the output of the models based on the SMC experience, resulting in an action. When the robot executes this action, it results in a movement of the arm. The system then observes the visual and proprioceptive consequences of that action and uses the observation of these new SMCs as training samples. In case of the forward model, the training data is $\{\mathbf{p}_{t+1}, \mathbf{x}_{t+1}, t\}$. The training data for the inverse model is $\{\mathbf{p}_t, \dot{\mathbf{x}}_t^{\mathbf{o}}, \mathbf{a}_t, t\}$, where $\dot{\mathbf{x}}_t^{\mathbf{o}} = \mathbf{x}_{t+1} - \mathbf{x}_t$ is the observed change in end effector.

2.2 Visual-shift model

The visual system provides information about the pose of the end effector and the target in visual coordinates. The visual-shift model applies a transformation, T, to these visual observations, so that the retinotopic (or camera) coordinate system is transformed into an internal visual coordinate system: $\mathbf{x} = T(\mathbf{x}^{\mathrm{r}})$ and $\tau = T(\tau^{\mathrm{r}})$, where $\mathbf{x}_t^{\mathrm{r}}$ is the pose of the end effector and τ_t^{r} the pose of the target, both in the retinotopic coordinate system.

The transformation T is updated based on the visually observed error in pointing, which is caused by the error of the model in predicting the effects of the applied action. T needs to counterbalance the visual transformation caused by the prism glasses. This can be done in different ways, but since in our experimental setup (see Section 3.1), the prism effect causes a rotation of visual observations, we chose to implement T as a rotation of the visual coordinates as well: $\mathbf{x} = T(\mathbf{x}^{\mathrm{r}}) = R^{\theta} \cdot \mathbf{x}^{\mathrm{r}}$, where the rotation

matrix R^{θ} applies a rotation over θ. The rotation angle θ is updated by the system at the end of each movement by $\theta_{t+1} = \theta_t + \eta \cdot \varepsilon$, where ϵ is the visually-observed terminal error. This error is based on the angular difference between the desired pose of the end effector at the end of the trajectory, τ_M, and the actual pose observed by the system after movement, $\mathbf{x}_{M+1}$: $\varepsilon = \angle\tau_M - \angle\mathbf{x}_{M+1}$, where M is the last time step in the action sequence. $\eta \in [0, 1]$ is the transformation learning rate, which determines the influence of the visual-shift model in the complete adaptation system.

2.3 Concurrent and terminal feedback

As in [15], we distinguish two different types of feedback: concurrent and terminal. In the first case, the end effector is continuously observed visually to obtain the internal pose $\mathbf{x}_t$, which is used by the inverse model to determine the action. This forms a visual closed-loop control system, where the effect of the action is visually observed and used in the next iteration. In the terminal-feedback condition, there is no visual closed-loop control, since the internal pose of the end effector is estimated based on proprioceptive information using the forward model. As a results, there will be a difference in terminal pointing error in the two conditions when SMCs are altered through prism exposure.

Another consequence of the feedback condition is the available data to train the inverse and forward model. In case of concurrent feedback, the new SMCs can reliably be observed over the complete trajectory. However, in the terminal condition, the SMCs related to the individual actions need to be estimated from the terminal feedback and will be incorrect due to the non-linear relations involved. We generate training data in the terminal-feedback condition by interpolating the internal visual pose of the end effector based on the visually observed end pose and the start pose estimated by the forward model.

3 Experiments

3.1 Simulation and model setup

We use a 2D simulated robotic setup to test our computational model of prism adaptation, see Figure 2a-c. The setup consists of a two degrees-of-freedom arm, a vision sensor observing the 2D position of the end effector, $\mathbf{x}_t = \{x_t^{\mathrm{x}}, x_t^{\mathrm{y}}\}$, and the target, $\tau_t = \{\tau_t^{\mathrm{x}}, \tau_t^{\mathrm{y}}\}$, and proprioceptive sensors in each of the joints, $\mathbf{p_t} = \{p_t^1, p_t^2\}$ giving information about the joint angles. In this setup, an action is a change in joint angles, $\mathbf{a}_t = \dot{\mathbf{q}}_t$. The visual observations are made from the bird's-eye perspective. In the experiment, pointing is done at two different heights, high and low, causing different arm poses during pointing. To experiment with the similarity of different poses, we add an arm-pose dimension to the input of the GPRs, with an associated length-scale parameter λ_{P}, which is by default set to 2.0 unless stated otherwise.

In a pointing trial, a target is positioned at a specific angle with respect to the robot. The system observes the target's position, $\mathbf{g}$, and then plans a target trajectory, $\{\tau_1, \ldots, \tau_M\}$, where $M = 5$ is the number of actions involved in the pointing trajectory (see Figure 2d). The prism glasses are implemented as a rotation of the visual coordinates, so that $\mathbf{y}^{\mathrm{r}} = R^{\gamma} \cdot \mathbf{y}^{\mathrm{t}}$, where $\mathbf{y}^{\mathrm{t}}$ is the true position, $\mathbf{y}^{\mathrm{r}}$ is the position in retinotopic

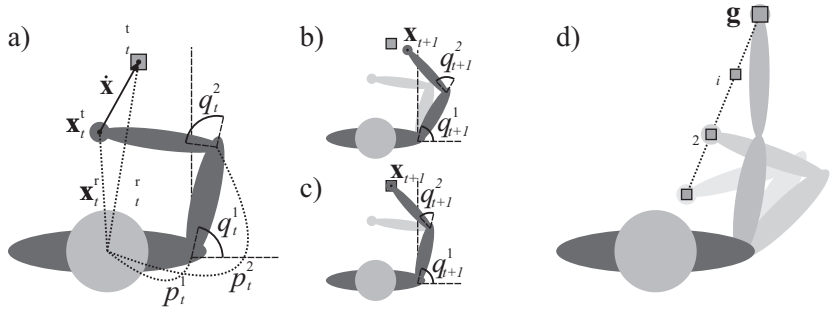

Fig. 2. Experimental setup. a) The robot has joint angles $\mathbf{q}_t$, the end effector is at true pose $\mathbf{x}_t^{\mathrm{t}}$, and the target at true pose τ_t^{t}. However, this information is not known to the system. Instead, the system observes (dotted lines) the proprioceptive signals $\mathbf{p_t}$, and the retinotopic coordinates of the end effector $\mathbf{x}_t^{\mathrm{r}}$ and the target τ_t^{r}. The model will suggest an action, which will change the joint angles. b) At first, the SMCs will not be correctly learned, and the action will result in an incorrect movement of the arm. c) After some time, the model will correctly estimate the action, so that $\mathbf{x}_{t+1} \equiv \tau_t$. d) Based on the visual target at position $\mathbf{g}$, the system plans a target trajectory giving a set of intermediate targets $\{\tau_1, \ldots, \tau_M\}$.

coordinates, and R^{γ} is a rotation matrix applying deviation angle γ. The application of the virtual prism glasses effects all visual observations, that is the observation of the end effector and the target.

We set the recency effect used in the forward and inverse models to $\lambda_{\mathrm{T}} = 1000$, and the transformation learning rate to $\eta = 0.35$. These parameters change the slopes of the learning curves and the influence of the forward and inverse models versus the visual-shift model. The general results are robust to small changes in these values.

3.2 Experimental setup

To compare the performance of our model to psychophysical data, we adopt the experiments performed by Redding and Wallace [15]. The task of the system is to point to visual targets. In the experiment, two different starting positions are used. In the proximal starting position, the end effector starts at the origin, and with a low arm pose. The distal starting position is halfway to the target, and with a high arm pose.

When the system has been initialized, as described below, we perform two sets of pre-tests, one for realignment and one for recalibration, to measure the performance before prism exposure. The prisms are then turned on in the simulation by rotating the visual observations over the origin with an angle of 11.4^{o}. Under prism exposure, the system performs 12 consecutive pointing movements from the distal starting position with either concurrent or terminal feedback. After that, the prisms are switched off, and the recalibration and realignment tests are performed again as post-tests. No adaptation learning is going on during the pre- and post-tests.

The realignment tests consist of a visual-shift test, a proprioceptive-shift test and a total-shift test, and are all done from the proximal starting position, i.e., different from the starting position during prism exposure. In the visual-shift test we measure subjective *straight-ahead*, by reading out $-\theta$, the negative value of the applied transformation

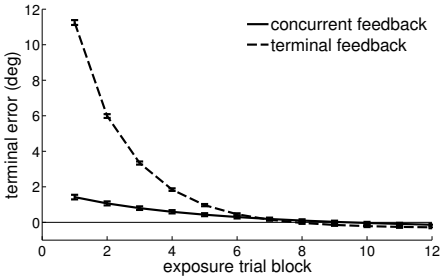

Fig. 3. The terminal pointing error as a function of pointing trial during prism exposure for concurrent and terminal feedback. The values are the means over 20 trials, and the error bars give the 95% confidence intervals.

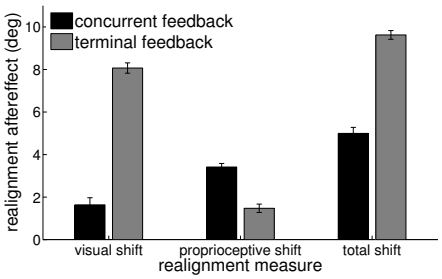

Fig. 4. The terminal error or aftereffect for the different realignment measures. With concurrent feedback, the visual shift remains small, while the proprioceptive shift is large, the reverse is true in case of terminal feedback.

in the visual-shift model. The proprioceptive-shift test measures proprioceptive straight ahead by having the system point without a visual target. To do so, the system sets a virtual target at 0° in internal coordinates. Finally, during the total-shift test, a visual target is presented at different angles $\{-20^{\circ}, -15^{\circ}, \ldots, 20^{\circ}\}$, and the system points to the targets. Pointing is always done without visual feedback.

The recalibration test is similar to the total-shift test, but the system is tested with the distal starting position, i.e., similar to the starting position during prism exposure. Moreover, we perform the test using different values for λ_{P}, to change the level of similarity in arm pose between distal and proximal starting conditions.

An experimental trial starts by initializing the system. The robot initiates 300 random movements in different parts of the work space and using the two different arm poses. Next, the system specializes on the pointing task by performing 24 pointing movements of $M = 5$ actions, using three different target positions (at -15°, 0° and 15°) and two different arm poses. Since this initialization is a noisy process and influences the performance of the system during the experiment, we repeat the experiment 20 times and report the mean values and 95% confidence intervals.

4 Results

Adaptation of pointing during prism exposure for the concurrent and terminal feedback condition is shown in Figure 3. In both conditions, the system quickly adapts to the prism effect, demonstrating accurate pointing behavior after 8 pointing trials. The error in the terminal condition is considerably larger than in the concurrent condition, which is expected, since the trajectory can be adjusted during pointing in the concurrent case. Negative error values can be observed in later pointing trials, showing overcompensation by the model, which is due to the collaboration of the forward and inverse models with the visual-shift model. These results correspond well with the psychophysical results observed by Redding and Wallace [15].

The size of the realignment aftereffect is shown in Figure 4. There is a clear difference in visual shift and proprioceptive shift between the concurrent and terminal

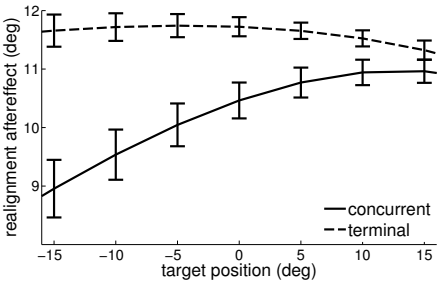

Fig. 5. The realignment aftereffect as a function of target position for concurrent and terminal feedback. The curves show different slopes depending on the feedback condition.

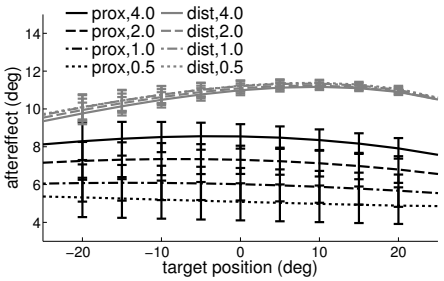

Fig. 6. The aftereffect for target positions tested at the proximal (realignment) and distal (recalibration) starting positions. λ_{P} indicates levels of pose similarity.

condition, matching the psychophysical results [15, 14]. The visual shift is small in the concurrent condition and large in the terminal condition. In the proprioceptive test, the aftereffects are reversed, with a large effect in the concurrent and a small effect in the terminal condition. The total aftereffects are approximately the sum of the visual- and proprioceptive aftereffects.

To investigate the generalization over space, we next investigate the total-shift (realignment) aftereffects as a function of target position for concurrent and terminal feedback (see Figure 5). With concurrent feedback, the aftereffect shows a positive slope towards the prismatic shift (11.4^{o}), whereas with terminal feedback, the curve shows a slight negative slope. This matches the psychophysical results [15].

Generalization over different actions is investigated by the aftereffects for the total shift and recalibration tests with proximal and distal starting position as shown in Figure 6. We use four different levels of similarity between the arm pose at those two starting positions, $\lambda_{\mathrm{P}} = \{4.0, 2.0, 1.0, 0.5\}$, where a higher value is a higher level of similarity. The results for the concurrent and terminal feedback condition are combined, as was done in [15]. The aftereffect is largest when the starting position is identical to the exposure phase, that is, the distal position. The curves for the distal position show a local generalization effect, with lower values for target positions that have not been used during prism exposure. This can be explained by the local learning in GPR used in forward and inverse model. The curves for the proximal position are more flat, indicating a more global generalization effect caused by the visual-shift model. Reducing the similarity in arm pose between exposure condition (distal) and the proximal condition during recalibration testing, i.e., lower values of λ_{P}, results in lower and more flat curves, indicating that aftereffects are weaker and mainly dominated by the global effect. These results correspond with psychophysical observations [15, 1].

5 Discussion

We presented a computational model capable of learning the sensorimotor contingencies involved in pointing, and of adapting to changes in these contingencies. Moreover, it accounts for the observations made in human prism-adaptation studies [14, 15, 1],

i.e. adaptation is achieved jointly through a specific local and a general global effect. Whereas local adaptation, or recalibration, is specific for the sensorimotor contingencies involved in the trained action, realignment, shows a global effect and generalizes to different actions

Using two feedback conditions, the realignment shows fundamentally different aftereffects for visual and proprioceptive tests. Where the visual shift is small and the proprioceptive shift is large for concurrent feedback, the reverse is true for terminal feedback. In our model, the visual shift is larger during terminal feedback because the observed pointing error at the end of the terminal feedback trials is larger compared to concurrent feedback trials. During concurrent feedback the end effector is continuously observed and the pointing error can be reduced while the movement is executed. In contrast, the results for the proprioceptive test are explained by the forward and inverse models. Since the sensory consequences of actions are directly observable in the concurrent condition, low-error training data is available which leads to quick adaptation of the inverse model and thus a large proprioceptive after effect. In the terminal condition, these training data need to be estimated based on the terminal observation, causing them to be less correct.

Our model gives a different explanation of prism adaptation compared to Redding and Wallace [15]. Where they consider recalibration to be a cognitive learning strategy and realignment to be an automatic process aligning different spatial maps, our model solely consists of automatic and low-level processes. Although cognitive strategies are an equally valid explanation, our model shows that similar effects can be reached with a simpler mechanism. Furthermore, in our case, the model does not include an explicit proprioceptive-shift model. Instead the results for the proprioceptive-shift test can be explained by adaptation of the forward and inverse models to changed SMCs. The presented model offers a potential system implementation for learning sensorimotor contingencies and emphasize the importance of the motor system for perception [8].

Currently, we assume a setup with a fixed eye-head system. However, in humans, both eye-head and head-hand systems are involved. The presented work is a first step towards modeling both systems. The current model is not able to learn multiple sensorimotor mappings. To account for dual adaptation effects observed in alternating prism-exposure experiments [18], future additions to the model will be necessary.

The model presented is directed at fostering our understanding of human adaptation as well as to improve on the current state of robotic systems. Inverse-kinematics learning with regression models has already been addressed by others, see e.g., [2, 9], but the addition of a recency effect and the synergy with the visual-shift model results in a fast adaptation to changes in sensorimotor contingencies.

Acknowledgments. This work was supported by the EU through the project eSMCs (FP7-IST-270212), and by the Swedish Foundation for Strategic Research.

Bibliography

[1] Baraduc, P., Wolpert, D.: Adaptation to a visuomotor shift depends on the starting posture. Journal of Neurophysiology 88(2), 973–981 (2002)

[2] D'Souza, A., Vijayakumar, S., Schaal, S.: Learning inverse kinematics. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (2001)

[3] Einhäuser, W., Martin, K., König, P.: Are switches in perception of the necker cube related to eye position? European Journal of Neuroscience 20(10), 2811–2818 (2004)

[4] Held, R., Hein, A.V.: Adaptation of disarranged hand-eye coordination contingent upon re-afferent stimulation. Perceptual and Motor Skills 8(3), 87–90 (1958)

[5] Held, R., Schlank, M.: Adaptation to disarranged eye-hand coördination in the distance-dimension. The American Journal of Psychology 72(4), 603–605 (1959)

[6] Kornheiser, A.: Adaptation to laterally displaced vision: A review. Psychological Bulletin 83(5), 783–816 (1976)

[7] Martin, T., Keating, J., Goodkin, H., Bastian, A., Thach, W.: Throwing while looking through prisms. ii. specificity and storage of multiple gaze-throw calibrations. Brain 119(4), 1199–1212 (1996)

[8] Nagel, S., Carl, C., Kringe, T., Märtin, R., König, P.: Beyond sensory substitution – learning the sixth sense. Journal of Neural Engineering 2, R13 (2005)

[9] Nguyen-Tuong, D., Peters, J.: Model learning for robot control: A survey. Cognitive Processing 12(4), 319–340 (2011)

[10] Nguyen-Tuong, D., Seeger, M., Peters, J.: Model learning with local gaussian process regression. Advanced Robotics 23(15), 2015–2034 (2009)

[11] O'Regan, J., Noe, A.: A sensorimotor account of vision and visual consciousness. Behavioral and Brain Sciences 24(5), 939–1031 (2001)

[12] Popović, M., Kootstra, G., Jørgensen, J.A., Kragic, D., Krüger, N.: Grasping unknown objects using an early cognitive vision system for general scene understanding. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 987–994. IEEE, San Francisco, CA, USA (2011)

[13] Rasmussen, C.E., Williams, C.: Gaussian Processes for Machine Learning. The MIT Press (2006)

[14] Redding, G.M., Wallace, B.: Components of prism adaptation in terminal and concurrent exposure: Organization of the eye-hand coordination loop. Perception and Psychophysics 44(1), 59–68 (1988)

[15] Redding, G.M., Wallace, B.: Generalization of prism adaptation. Journal of Experimental Psychology: Human Perception and Performance 32(4), 1006–1022 (2006)

[16] Redding, G.M., Wallace, B.: Intermanual transfer of prism adaptation. Journal of Motor Behavior 40(3), 246–262 (2008)

[17] Schaal, S., Atkeson, C.G., Vijayakumar, S.: Scalable techniques from nonparametric statistics for real-time robot learning. Applied Intelligence 17(1), 49–60 (2002)

[18] Welch, R., Bridgeman, B., Anand, S., Browman, K.: Alternating prism exposure causes dual adaptation and generalization to a novel displacement. Perceptual Psychophysics 54(2), 195–205 (1993)