

Humanoid Oculomotor Control Based on Concepts of Computational Neuroscience

Tomohiro Shibata^{*1}, Sethu Vijayakumar^{*2}, Jörg Conradt^{*3}, and Stefan Schaal^{*2}

^{*1} Kawato Dynamic Brain Project, ERATO, JST

^{*2} Computer Science & Neuroscience, University of Southern California

^{*3} Institute of Neuroinformatics, University/ETH Zurich

Keywords: oculomotor control, computational neuroscience, learning control

Abstract

Oculomotor control in a humanoid robot faces similar problems as biological oculomotor systems, i.e., the stabilization of gaze in face of unknown perturbations of the body, selective attention, stereo vision, and dealing with large information processing delays. In this paper we suggest control circuits to realize three of the most basic oculomotor behaviors, the vestibulo-ocular and optokinetic reflex, smooth pursuit, and saccades for overt visual attention. Each of these behaviors was derived from inspirations from computational neuroscience, which proves to be a viable strategy to explore novel control mechanisms for humanoid robotics. Our implementations on a humanoid robot demonstrate good performance of the oculomotor behaviors that appears natural and human-like.

1 Introduction

1.1 Research Objectives

The goal of our research is to understand the principles of information processing of the human brain, with a focus on basic sensorimotor control and the hope to expand this scope increasingly more towards more cognitive topics. As a research strategy, we chose an approach that emphasizes the interplay between computational neuroscience and humanoid robotics. In this approach, research topics are initially investigated from the present stage of knowledge of neurobiology and, subsequently, abstract computational models are created that can be implemented on a humanoid robot to accomplish interesting behavioral goals. Control theory and learning theory are employed to examine the validity of the models and the success of the models in actual behavior is investi-

gated. Theoretical and experimental insights are then used to re-evaluate biological data and the present stage of modeling, which usually leads to suggestions for improvement in both neuroscientific research and computational modeling. For the purpose of this research strategy, we developed a humanoid robot system with 30 degree-of-freedom (DOFs), each of which is hydraulically operated and mimics the compliance of humans by means of impedance control in each joint. Kinematics and dynamics of the robot are as close as possible to the human counterpart.

In this paper, we present results of our research in the field of oculomotor control, i.e., learning the vestibulo-ocular and optokinetic reflex (VOR-OKR), smooth pursuit, and saccades for overt visual attention. Oculomotor control is one of the best investigated areas in computational neuroscience due to the relative simplicity of the primate oculomotor systems and rich set of oculomotor behaviors, including reflexes and adaptation. For example, the monkey's oculomotor system can be approximated by a second order linear system. It has only 3 DOFs per eye, and usually only 1 DOF is used in neurobiological experiments. Furthermore, in the case of reflex behaviors, experimenters can eliminate the effects of attention, and easily elicit the reflex behavior by simple stimuli such as rotation of the head, spot lights, random dot displays, etc. Such simple experimental paradigms are ideally suited to draw comparisons between biological knowledge, computational models, and empirical evaluations in robotic experiments.

1.2 Robotic Head Setup

In the next sections, we will present computational models for the three oculomotor behaviors we examined and the corresponding experimental results. In all experiments, the same robotic platform was used,




Figure 1: Our humanoid head

a robotic head of a Humanoid Robot [1].

The robot head has 7 DOFs in total, a neck with 3 DOFs and two camera eyes, each equipped with 2 independent DOFs, arranged as pan and tilt. In order to provide high-resolution vision simultaneously with large-field peripheral vision, the robot employs two cameras per eye, a foveal camera (24 deg view-angle horizontally) and a wide-angle camera (100 deg view-angle horizontally). This setup mimics the foveated retinal structure of primates, and it is also essential for an artificial vision system in order to obtain high resolution vision of objects of interest while still being able to perceive events in the peripheral environment.

The learning controller is implemented with the real-time operating system VxWorks using several parallel Motorola PowerPC processors in a VME rack. Visual processing is performed out of specialized hardware, a Fujitsu tracking vision board and QuickMag (Japan) color vision tracking system. The Fujitsu tracking vision board calculates retinal slip and retinal slip velocity information of each eye at 30Hz. The QuickMag system returns the centroid of blobs of pixels of pre-specified colors in the environment. Up to six different colors can be tracked simultaneously at 60Hz sampling rate.

The oculomotor control loop runs at 420 Hz, while the vision control loop runs at 60 Hz due to the QuickMag video processing rate. The oculomotor control loop implements a strong spring and damping term such that the nonlinearities of the oculomotor system due to hydraulics and attached cables become negligible.

2 Vestibulo-Ocular-Reflex (VOR)

This section outlines the computational model of VOR-OKR we developed [1], shown schematically in Fig. 2. The VOR reflex serves to keep the eyes fixed on a target in case that there is a mechanical perturbation of the head, e.g., as caused by locomotion. The OKR has a similar functionality, just that it is triggered by a movement of the entire visual field, which it tries to compensate for. A typical movement would be elicited in a movie theater when the entire scene on the screen moves. The inputs to the VOR-OKR system are i) the visual target in camera coordinates and ii) an angular velocity signal generated from a gyroscopic sensor due to perturbations of the robot's body; since the sensor is attached to the head, the signal is referred to as "head angular velocity". From the target position and eye position, retinal slip error and retinal slip velocity can be computed. In the simplest case, the ideal compensatory desired movement of the eyes would be the negative of the retinal slip error velocity, but, in general, a nonlinear transformation from retinal slip error velocity to eye movement is needed due to off-axis effect, i.e. the fact that head axis and eye axis are not collinear. The retinal error signals are also used as input to a PD controller in the bottom part of Fig. 2. The gains of this PD controller have to be kept rather small due to the delays incurred in visual information processing. The output of the PD controller serves primarily as a teaching signal to the feedback error learning system. However, it is also needed to stabilize the crude feedforward controller in the shaded block of Figure 2. Without the feedback input, the feedforward controller would only be marginally stable due to the floating integrator. As described in a later section, eligibility traces, realized by a second order filter, are also used as inputs to the learning system.

The entire control systems is quite similar to what has been discovered in the primate oculomotor system. The a priori existing *feedforward controller* provides some crude functionality of the VOR while the existing *feedback controller* provides acceptable OKR performance for slowly changing visual targets and acts as a compensatory negative feedback for the VOR module. These systems form what is called the "direct pathway" of oculomotor control in biology. By adding a learning controller trained with the FEL strategy in the indirect pathway (see Fig. 2), excellent VOR performance can be accomplished even if the feedback pathway has large delays. Furthermore, the OKR performance is improved to some extent. This learning network, known to be located in the primate cere-

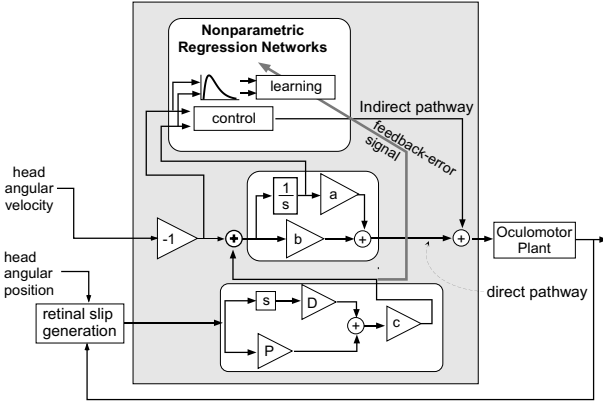


Figure 2: Our VOR-OKR model

bellum, acquires an inverse dynamics model of the oculomotor plant during the course of learning that compensates for the missing performance of the crude feedforward controller in the shaded box (c.f. Fig. 2). The coordination of a direct and indirect pathway is analogous to how the cerebellar pathway acts in parallel to the brainstem pathways [2]. As discussed in [1], this control system is equally suitable for both biological and robotic oculomotor control. However, for the robotic control, we augment the biologically plausible model with a fast nonlinear learning scheme based on nonparametric regression networks [3].

2.1 Learning with the delayed-error signal

For successful feedback-error-learning, the time-alignment between input signals and the feedback-error signal is theoretically crucial, and, thus, additional techniques are required in the case of delayed sensory feedback. For instance, if a perturbation of the head or body has frequency components that are much faster than the delay in the feedback pathway during VOR learning, the phase delay in the feedback pathway gets large such that learning speed decreases, or learning can even become unstable in the worst case.

To solve this “temporal credit assignment problem”, the concept of eligibility traces has been suggested in both biological modeling and machine learning [4]. For neurons in the brain, it is assumed that a second messenger would tag a synapse as eligible for modification. This “tag” would decay with an appropriate time constant, thus forming a temporal eligibility window. Schweighofer et al. proposed a biologically plausible learning model for saccadic eye movement, and modeled the second messenger as a second order linear filter of the input signals to the learning




Figure 3: Time course of the rectified mean retinal slip: While the dashed line represents data obtained without eligibility trace, the solid line shows data acquired with eligibility trace.

system[5]. Using a second order filter is important since, in contrast to a first order filter, the impulse response has a unimodal peak at a delay time determined by the time constant of the filter. For successful learning, the delay time only has to roughly coincide with the actual delay of the sensory feedback. In related work, Fagg proposed a cerebellar learning model where eligibility traces modeled by such a second order filter are employed[6]. Applying this technique to FEL, we complete our final learning control system (see Fig. 2) where the impulse response of a second order linear filter is put just before the “learning” box. Results of experimental evaluations on our anthropomorphic robot based on this proposed control system will be discussed in the next section.

2.2 Experimental result

Head movement was generated by three superimposed sinusoidal signals with frequencies of 0.6, 2.0, and 3.0 Hz and amplitude of 0.1 rad, respectively. A frequency of 3.0 Hz is high enough to obtain sustained blurred visual images. Fig. 3 shows the time course of the rectified retinal slip during learning obtained from a moving average using a one second time window. While the dashed line represents data obtained from learning without eligibility traces, the solid line shows data acquired with eligibility traces. This figure shows that eligibility traces are necessary for successful learning as the the retinal slip does not decrease without using the traces. It should also be noted that learning

proceeds quite rapidly such that in less than a half a minute, the initial errors are reduced by a factor of two. Longer learning results in a further reduction of the error [1].

3 Smooth Pursuit

Smooth pursuit refers to the oculomotor behavior of tracking a moving target as accurately as possible; for instance such behaviors are needed to visually inspect moving object. Strong evidences (e.g.,[7]) support the idea that smooth pursuit employs some predictive control mechanism. From the control theoretical view, this observation is not surprising since smooth pursuit cannot be achieved by a simple visual negative feedback controller due to the long delays inherent in visual information processing (e.g. around 70 ms in our humanoid vision system, and around 50 ms in the human brain).

In the field of robot vision, many projects investigated visual servoing, but, to our knowledge, without ever examining a smooth pursuit controller that has similar features and performance as that in primates. One of the most related pieces of research is in Bradshaw et al. [8], which employed a Kalman filter for prediction. However, these authors assumed prior knowledge of the target dynamics and, thus, avoided to address how unknown target motion can be tracked accurately. In contrast, in this paper, we present a biologically inspired smooth pursuit controller that learns to predict the visual target velocity in head coordinates based on fast on-line statistical learning of the target dynamics. In the following sections, we will first explain the setup of our smooth pursuit model, explain the learning component, and, in Section 3.2, describe how to accelerate learning speed. The suggested control system is evaluated in simulations and experiments with our humanoid robot.

3.1 Cascade of predictor and controller

Fig. 4 presents our smooth pursuit model. It consists of two subsystems: one is a target velocity predictor, and the other is an inverse model controller of the oculomotor system. In this control diagram, s and $1/s$ are Laplace transform operators denoting differentiation and integration, respectively. Δ stands for a constant delay element. e , \dot{e} , E , and \dot{E} are the retinal error and the retinal slip, the eye angular position, and the eye angular velocity, respectively.

The predictor outputs an estimate of the current target velocity $\hat{\mathbf{T}}(t)$ out of a history of past estimated

target angular positions $T(t - \Delta)$ and velocities $\dot{T}(t - \Delta)$. In linear systems, the state predictor of a n -th order linear system can be defined as:

$$\mathbf{x}_{t+1} = A\mathbf{x}_t \quad (1)$$

where \mathbf{x} is the $n \times 1$ state vector and A is the $n \times n$ state transition matrix. As we are only interested in velocity prediction in this paper, we reduce equation 1 to focus only on the the states that are identified with target velocities, not positions:

$$\hat{\mathbf{T}}_{t+1} = \mathbf{A}_2 \mathbf{x}_t \quad (2)$$

where \mathbf{A}_2 is the appropriate submatrix of A corresponding to the target velocity component.

Since, similar to the primates' eye plant, many artificial oculomotor systems can be approximated by a second order linear dynamical system, a linear inverse dynamics model suffices for feedforward control of the eye system. It receives as input the desired eye angular position, velocity, and acceleration. The feedforward controller is stabilized by a PD feedback pathway whose gains need to be rather low due to its operating out of strongly delayed signals.




Figure 4: Our smooth pursuit model

3.2 Learning the discrete predictor from the delayed signals

The learning scheme in Fig. 4 may appear difficult to implement as it has to learn the target dynamics out of the history of past estimated target states and the delayed retinal error. As will become apparent in the next paragraphs, however, a straightforward development allows us to solve this learning problem.

At time t , the predictor can only see the delayed estimated target state $\mathbf{x}_{t-\Delta}$. The corresponding discrete target velocity prediction is represented as

$$\hat{\mathbf{T}}_t = f(\mathbf{x}_{t-\Delta}, \mathbf{w}_t) \quad (3)$$

where \mathbf{w} is a parameter vector. Let $\dot{\xi}$, the velocity prediction error, equal $\dot{T} - \hat{\dot{T}}$, and let the loss function

J be the simple squared error:

$$J = \frac{1}{2} \dot{\xi}_t^2 \quad (4)$$

Thus, a gradient descent learning rule for w can be written as:

$$\left(\frac{dw_i}{dt}\right)_t = -\epsilon \left(\frac{\partial J}{\partial w_i}\right)_t = \epsilon \left(\frac{\partial f}{\partial w_i}\right)_{t-\Delta} \dot{\xi}_t \quad (5)$$

with ϵ denoting the learning rate. If we can make the assumption that the predicted target velocity \hat{T} will be tracked accurately by the robot without delay, we can regard the retinal slip as the prediction error given by $\dot{\xi} = \dot{T} - \hat{T} \simeq \dot{T} - \dot{E} = \dot{e}$. The learning rule can thus be rewritten as:

$$\left(\frac{dw_i}{dt}\right)_t = \epsilon \left(\frac{\partial f}{\partial w_i}\right)_{t-\Delta} \dot{e}_t \quad (6)$$

Note that the time alignment of the predictor output f and the error $\dot{\xi}$ ($\simeq \dot{e}$) needs to be correct for successful minimization of the loss function J . Since the predictor has no access to \dot{e}_t at time t , a modified learning rule is required by introducing a delayed form of Eq. 6:

$$\left(\frac{dw_i}{dt}\right)_t = \epsilon \left(\frac{\partial f}{\partial w_i}\right)_{t-2\Delta} \dot{e}_{t-\Delta} \quad (7)$$

Thus, the predictor is required to keep the information $\partial f / \partial w_i$ in memory for the duration of Δ .

In summary, it is important to use the most recent information for prediction, but to use one delayed by Δ for learning in order to achieve successful learning and control.

3.3 Experimental result

We implemented the on-line learning of smooth pursuit on our humanoid robot. Our model does not rely on any specific or particular learning method. For tracking an target having a simple linear dynamics such as a swinging pendulum, we employed an adapted version of recursive least squares (RLS), and obtained excellent results. For more complex target motion with nonlinear dynamics, we replaced RLS in our smooth pursuit controller with nonparametric regression networks we developed. This system was used to learn a periodic motion generated by van der Pol equations implemented on a separate industrial robot in our laboratory. Fig. 5 shows the excellent learning results of this experiment.




Figure 5: Smooth pursuit of a target following a trajectory that is generated by Van der Pole equation. The upper figure shows the time course of the angular position of the visual target (dotted) and the eye (solid). The lower figure presents the time course of the rectified mean retinal error (smoothed with moving average of time window 1 s)

4 Saccade for Overt Visual Attention

Visual attention involves directing a “spotlight” of attention [9] to interesting areas, extracted from a multitude of sensory inputs. Most commonly, attention will require to move the body, head, eyes, or a combination of these in order to acquire the target of interest with high-resolution foveal vision, referred to as ‘overt’ attention, as opposed to covert attention which does not involve movement.

There has been extensive work in modeling attention and understanding the neurobiological mechanisms of generating the visual “spotlight” of attention [10], both from a top-down[11] and a bottom-up perspective [12, 13] - albeit mainly for static images. From the perspective of overt shift of foci, there has been some work on saccadic eye motion generation using spatial filters [14], saccadic motor planning by integrating visual information [15], social robotics [16], and humanoid robotics [17]. In contrast to this previous work, our research focus lies on creating a biologically inspired approach to visual attention and oculomotor control by employing theoretically sound computational elements that were derived from models of cortical neural networks, and that can serve for comparisons with biological behavior. We also emphasize real-time performance and the integration of the attention system on a full-body humanoid robot that is

not stationary in world coordinates. As will be shown below, these features require additional computational consideration such as the remapping of a saliency map for attention after body movement. In the following sections, we will first give an overview of the attentional system's modules, then explain the computational principles of each module, before we provide some experimental evaluations on our humanoid robot.

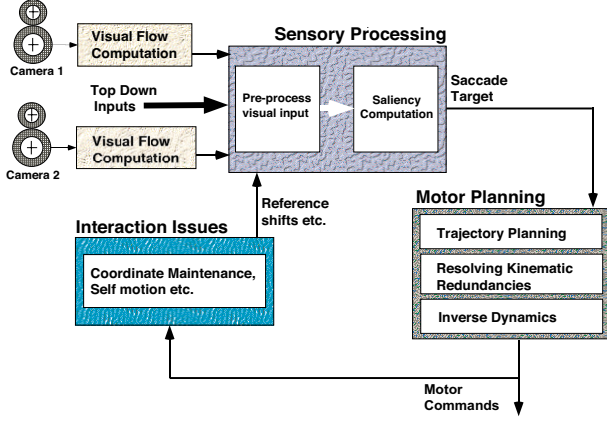


Figure 6: A schematic block diagram of the various modules involved in the system for implementing overt visual attention

4.1 Sensor pre-processing and integration

The key element of our Sensory Pre-Processing block (Fig. 6) is a competitive dynamical neural network, derived in Amari and Arbib's [18] *neural fields* approach for modeling cortical information processing. The goal of this network is to take as input spatially localized stimuli, have them compete to become the next saccade target, and finally output the winning target. For this purpose, the sensory input pre-

processing stage takes the raw visual flow $V_F(\mathbf{x}, t)$ as inputs to the *stimulus dynamics*, a first order dynamical system. Using \mathbf{x} to denote the position of a stimulus in camera coordinates, the stimulus dynamics is:

$$\dot{S}(\mathbf{x}) = -\alpha S(\mathbf{x}) + VisInp(\mathbf{x}, t) \quad (8)$$

where

$$VisInp(\mathbf{x}, t) = \int_R G(\mathbf{x}, \mathbf{t}) * \exp(-\mathbf{x}^2/2\sigma^2) d\mathbf{x} \quad (9)$$

$$G(\mathbf{x}, t) = V_F(\mathbf{x}, t) + \gamma * [\dot{V}_F(\mathbf{x}, t)]_+ \quad (10)$$

Eq.(10) enhances the raw visual flow vector when it is increasing to emphasize new stimuli in the scene, while Eq.(9) implements a Gaussian spatial smoother of the stimuli to reduce the effects of noise. The variable α was set to a value of 100 in our experiments. The top of Fig. 7a shows an example of a typical stimulus pattern in the two dimensional neural network due to a moving object at the top-left of the camera image. In general, we could have multimodal sensory inputs, e.g. from color detectors, edge detectors, audio input, etc., feeding into Eq.(10) as a sensory signal. As suggested by Niebur, Itti and Koch [12, 13], it would be useful to weight these inputs according to their importance in the scene, usually based on some top-down feedback or task-specific biasing (e.g., if we know that color is more important than motion).

This stimulus dynamics feeds into a *saliency map* [9], essentially a winner-take-all (WTA) network which decides on a winning stimulus from many simultaneous stimuli in the camera field. The winning stimulus will become the next saccade target or focus of overt attention. The WTA network is realized based on the theory of *neural fields*, a spatial neural network inspired by the dynamics of short range excitatory and long range inhibitory interactions in the neo-cortex [18, 19]. The activation dynamics $u(\mathbf{x}, t)$ of the saliency map is expressed as:

$$\tau \dot{u}(\mathbf{x}) = -u(\mathbf{x}) + S(\mathbf{x}) + \mathbf{h} + \sum_{\mathbf{x}'} \mathbf{w}(\mathbf{x}, \mathbf{x}') \sigma(u(\mathbf{x}')) \quad (11)$$

Here, h is the base line activation level within the field, $S(\mathbf{x}, t)$ is the external stimulus input (Eq.8), $\mathbf{w}(\mathbf{x}, \mathbf{x}')$ describes the coupling strength between all the units of the network, and $\sigma(u)$ controls the local threshold of activation. Depending on the choice of parameter h and the form of σ and \mathbf{w} , the activation dynamics of Eq.(11) can have various stable equilibrium points [18]. We are interested in a solution which has uniform activation at base line level in the absence of external

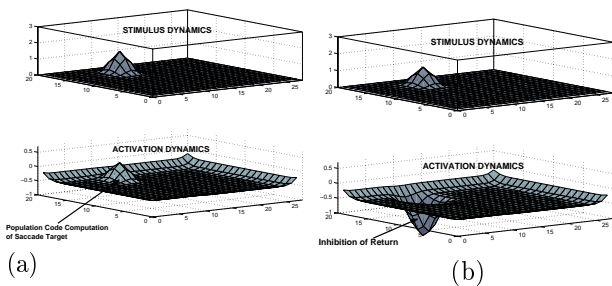


Figure 7: A snap shot of the stimulus and activation dynamics just (a) before and (b) after the saccade

stimuli, and which forms a unimodal activation pattern at the most significant stimulus in the presence of stimuli that are possibly dispersed throughout the spatial network. This is achieved by choosing a transfer function:

$$\sigma(u) = 1/(e^{(-cu)} + 1) \quad (12)$$

with constant $c \gg 1$ and an interaction kernel with short range excitation and long-range inhibition term H_0 :

$$\mathbf{w}(\mathbf{x}, \mathbf{x}') = k e^{-(\mathbf{x} - \mathbf{x}')^2 / \sigma_w^2} - \mathbf{H}_0 \quad (13)$$

The constants were fixed at $\tau = 0.01$, $h = -0.5$, $H_0 = 0.75$, $k = 4$, $\sigma_w^2 = 1.4$, and $c = 5000$, the values of which were decided based the magnitude of the stimulus dynamics $S(\mathbf{x}, \mathbf{t})$, as outlined in [18].

In addition to the stimulus driven dynamics, we also suppress the activation of the most recently attended location by adding a large negative activation in Eq.(10) at the location of the last saccade target. This strategy implements an *inhibition of return* [13] and ensures that the robot does not keep attending to the same location in the continuous presence of an interesting stimuli. The plots at the bottom of Fig. 7(a)(b) illustrate the behavior of the activation dynamics just before and after an attention shift, including the effect of the negative activation after the saccade.

4.2 Planning and generation of motor commands

Given a new saccade target, extracted from the saliency map, the direction of gaze needs to be shifted to the center of this target. Since fifth order splines are a good approximation of biological movement trajectories [20, 21], we use this model to compute a desired trajectory from the current position \mathbf{x}_0 to the target \mathbf{x}_f , all expressed in camera coordinates.

The camera-space trajectory is converted to joint space by inverse kinematics computations based on Resolved Motion Rate Control (RMRC) [22]. We assume that only head and eye motion is needed to shift the gaze to the visual target, an assumption that is justified given that the target was already visible in the peripheral field of view. For the time being, the inverse kinematics computation is performed for the right eye only, while the left eye performs exactly the same motion as the right eye. Thus, we need to map from a 2D camera space of the right eye to a 5D joint space, comprised of pan and tilt of the camera, and 3 DOFs of the robot's neck. To obtain a unique inverse,

we employ Liegeois [22] pseudo-inverse with optimization:

$$\dot{\boldsymbol{\theta}} = \mathbf{J}^\# \dot{\mathbf{x}} + (\mathbf{I} - \mathbf{J}\mathbf{J}^\#) \mathbf{k}_{\text{null}} \quad (14)$$

where $\mathbf{J}^\# = \mathbf{J}^T (\mathbf{J}\mathbf{J}^T)^{-1}$

\mathbf{k}_{null} is the gradient of an optimization criterion w.r.t. the joint angles $\boldsymbol{\theta}$. The second term of the Eq.(14) is the part that controls the movement in the null space of the head-eye system. Any contribution to $\dot{\boldsymbol{\theta}}$ from this term will not change the direction of gaze but will only change how much we use the head or eye DOFs to realize that gaze. As optimization criterion we chose:

$$L = \frac{1}{2} \sum_i w_i (\theta_i - \theta_{def,i})^2 \quad (15)$$

resulting in

$$k_{\text{null},i} = \frac{\partial L}{\partial \theta_i} = w_i (\theta_i - \theta_{def,i}) \quad (16)$$

$$(17)$$

This criterion keeps the redundant DOFs as close as possible to a default posture $\boldsymbol{\theta}_{def}$. Adding the weights w_i allows giving more or less importance to enforcing the optimization criterion for certain DOF—this feature is useful to create natural looking head-eye coordination.

Once the desired trajectory is converted to joint space, it is tracked by an inverse dynamics controller using a learned inverse dynamics model [23].

4.3 Experimental result

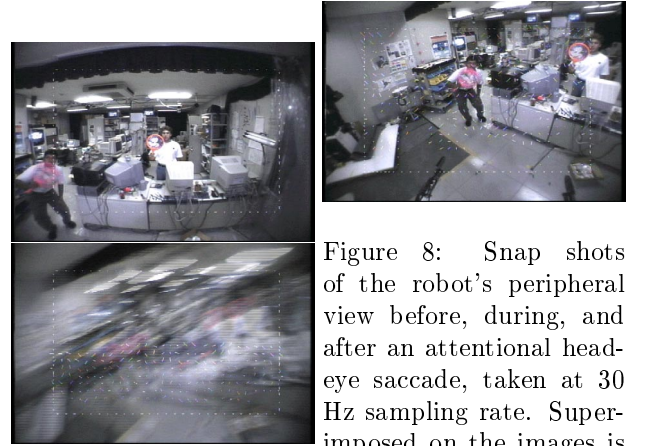


Figure 8: Snap shots of the robot's peripheral view before, during, and after an attentional head-eye saccade, taken at 30 Hz sampling rate. Superimposed on the images is the visual flow field.

We implemented the visual attention system on our humanoid robot. The stimulus dynamics and saliency map had 44x44 nodes, i.e., twice the length and width

of the 22x22 nodes of the visual flow grid of the peripheral vision. This extended size assured that after a saccade, the remapping of the saliency map and stimulus dynamics could maintain stimuli outside of the peripheral vision for some time. The Jacobian needed for the inverse kinematics computation was estimated with linear regression from data collected from moving the head-eye system on randomized periodic trajectories for a few minutes. Due to the limited range of motion of the eye and head DOFs, the Jacobian could be assumed to be constant throughout the entire range of motion of head-eye system, which was confirmed by the excellent coefficient of determination of the regression of the Jacobian. The saliency map was able to determine winning targets at about 10Hz, which is comparable to the capabilities of the human attentional system.

An illustration of the working of the attentional system is provided in Fig. 8. The top image shows the robot's right eye peripheral view of the lab, focusing on the person in the middle of the image. At the bottom left part of the image, another person was waving a racket to attract the robot's attention. This motion elicited a saccade, recognizable from the middle image of Fig. 8 which shows the visual blur that the robot experienced during the movement. The bottom image of Fig. 8 demonstrates that after the saccade, the robot was correctly focusing on the new target. Note that the three images were sampled at 30Hz, indicating that the robot performed a very fast head-eye saccade of about 60ms duration, which is again comparable to human performance. Our future work will augment the attentional system with more sensory modalities, including the learning of weighting of sensory modalities for different tasks.

5 Conclusion

In this paper, we presented our research on humanoid oculomotor control, focusing on models of the VOR-OKR reflex system, smooth pursuit, and saccade generation for visual attention based on concepts of computational neuroscience. We demonstrated the performance of these oculomotor behaviors by using our humanoid robot head that is part of our Humanoid Robot, developed for the purpose of brain science research. In all given examples, the robot control mechanisms were derived from principles of computational neuroscience that proved to be generate viable solutions for robotic control with good performance.

Currently our work is dedicated to the integration of the VOR-OKR, smooth pursuit, and saccade sys-

tem to give our humanoid the basis of the visual perception. For this purpose, all behaviors need to send their motor commands to a final common pathway, and it needs to be ensured that the behaviors do not interfere negatively. These integration issues are surprisingly complex and will require parallel investigations of the primate counterpart.

References

- [1] T. Shibata and S. Schaal. Biomimetic gaze stabilization based on feedback-error-learning with non-parametric regression networks. *Neural Networks*, 14(2):201–216, 2001.
- [2] H. Gomi and M. Kawato. Adaptive feedback control models of the vestibulo-cerebellum and spinocerebellum. *Biological Cybernetics*, 68:105–114, 1992.
- [3] S. Vijayakumar and S. Schaal. Fast and efficient incremental learning for high dimensional movement systems. In *IEEE Int. Conf. on Robot. and Automat.*, pages 1894–1899, 2000.
- [4] A.G. Barto, R.S. Sutton, and C.W. Anderson. Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE Trans. on Systems, Man, and Cybern.*, SMC-13:834–846, 1983.
- [5] N. Schweighofer, M.A. Arbib, and P.F. Dominey. A model of the cerebellum in adaptive control of saccadic gain. *Biol Cybern.*, 75:19–28, 1996.
- [6] A.H. Fagg, N. Sitkoff, A.G. Barto, and J.C. Houk. Cerebellar learning for control of a two-link arm in muscle space. In *Proc. of IEEE Conf. on Robotics and Automation*, pages 2638–2644, 1997.
- [7] S.G. Whittaker and G. Eaholtz. Learning patterns of eye motion for foveal pursuit. *Assoc Res Vis Ophthal*, 23(3):393–397, 1982.
- [8] K.J. Bradshaw, I.D. Reid, and D.W. Murray. The Active Recovery of 3D Motion Trajectories and Their Use in Prediction. *IEEE Trans. on Pattern Analysis and Machine Intel.*, 19(3), 1997.
- [9] C. Koch and S. Ullman. Selecting one among the many: A simple network implementing shifts in selective visual attention. A.I. Memo 770, Massachusetts Institute of Technology, January 1984.
- [10] E. Neibur and C. Koch. Computational architectures for attention. In R. Parasuraman, editor, *The Attentive Brain*, pages 163–186. MIT Press, Cambridge, MA, 1998.
- [11] R. Parasuraman. *The Attentive Brain*. MIT Press, Cambridge, MA, 1998.
- [12] L. Itti and C. Koch. A comparison of feature combination strategies for saliency-based visual attention systems. In *Proc. SPIE Human Vision and Electronic Imaging IV (HVEI'99)*, San Jose, CA, volume 3644, pages 473–82, January 1999.

- [13] L. Itti and C. Koch. A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research*, 40(10-12):1489–1506, 2000.
- [14] R. Rao and D. Ballard. Learning saccadic eye movements using multiscale spatial filters. In G. Tesauro, D. S. Touretzky, and T. K. Leen, editors, *Advances in Neural Information Processing Systems 7 (NIPS 94)*, pages 893–900. MIT Press, 1995.
- [15] K. Kopecz and G. Schoner. Saccadic motor planning by integrating visual information and pre-information on neural dynamic fields. *Biological Cybernetics*, 73:49–60, 1995.
- [16] C. Breazeal and B. Scassellati. A context dependent attention system for a humanoid robot. In *Proc of IJCAI-99*, pages 1146–1151, 1999.
- [17] J.A. Driscoll, R.A. Peters II, and K.R. Cave. A visual attention network for a humanoid robot. In *Proc Int Conf Intell Robot Syst (IROS-98)*, pages 1968–1974, 1998.
- [18] S. Amari. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27:77–87, 1977.
- [19] S. Amari and M.A. Arbib. Competition and cooperation in neural nets. In Metzler, editor, *Systems neuroscience*, pages 119 – 165. Academic Press, New York, 1977.
- [20] M. Kawato. Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9:718–727, 1999.
- [21] G.R. Barnes. Visual-vestibular interaction in the control of head and eye movement: The role of visual feedback and predictive mechanisms. *Progress in Neurobiology*, 41:435–472, 1993.
- [22] A. Liegeois. Automatic supervisory control of the configuration and behavior of multibody mechanisms. *IEEE Transactions on Systems, Man, and Cybernetics*, 7:868–871, 1977.
- [23] S. Vijayakumar and S. Schaal. LWPR : An $O(n)$ algorithm for incremental real time learning in high dimensional space. In *Proc. International Conference on Machine Learning (ICML 2000)*, pages 1079–1086, 2000.