

Visuo-inertial stabilization in space-variant binocular systems [☆]

Francesco Panerai ^{a,*}, Giorgio Metta ^b, Giulio Sandini ^b

^a *Laboratoire de Physiologie de la Perception et de l'Action (LPPA), Collège de France, 11 place Marcelin Berthelot, 75005 Paris, France*

^b *Laboratory for Integrated Advanced Robotics (LIRA), Department of Communication, Computers and Systems Science, University of Genoa, Via Opera Pia 13, 16145 Genoa, Italy*

Abstract

Stabilization of gaze is a major functional prerequisite for robots exploring the environment. The main reason for a “steady-image” requirement is to prevent the robot’s own motion to compromise its “visual functions”. In this paper we present an artificial system, the LIRA robot head, capable of controlling its cameras/eyes to stabilize gaze. The system features a stabilization mechanism relying on principles exploited by natural systems: an *inertial* sensory apparatus and images of *space-variant* resolution. The inertial device measures angular velocities and linear acceleration along the vertical and horizontal fronto-parallel axes. The space-variant image geometry facilitates real-time computation of optic flow and the extraction of first-order motion parameters. Experiments which describe the performance of the LIRA robot head are presented. The results show that the stabilization mechanism improves the reactivity of the system to changes occurring suddenly at new spotted locations. ©2000 Elsevier Science B.V. All rights reserved.

Keywords: Inertial sensors; Image stabilization; Visuo-inertial integration; Space-variant binocular vision

1. Introduction

Stabilization mechanisms have developed in natural systems as necessary functional elements to enable consistent treatment of visual information. These mechanisms are indeed important when the visual system needs to exploit its maximum acuity, estimate self-motion or discriminate motion of the visual scene. As a matter of fact, images falling on the retinas are required to be relatively still, particularly in all those tasks requiring the brain’s neural circuitry to discrim-

inate/extract/recognize relevant visual features. In humans, for example, it is well known that the stabilization of the whole visual field enhances sensitivity to individual moving objects and to motion parallax [27], contributing to enhancement in the perception of distance, form [35], and relative depth [12]. In robotics too, image stabilization mechanisms should be considered important functional elements. Robots exploring the environment will certainly benefit from a stable gaze-line. Two cases are of particular interest: (i) when performing dynamic visual measurements during their own movement [33]; (ii) when using fixation as a reference while navigating or manipulating [3]. From an implementation perspective, how should one go about achieving robust performance and efficiency at the same time? What is the sensory information that should be considered to achieve the goal? Does the sensing process satisfy the dynamics of the stimuli?

[☆] The research described in this paper has been supported by the Italian Space Agency (ASI) and the EU Esprit project NARVAL.

* Corresponding author. Tel.: +33-1-4427-1389; fax: +33-1-4427-1382.

E-mail address: francesco.panerai@college-de-france.fr (F. Panerai).

What are the computational cost of a particular implementation and the resources available to the robot?

In animals with fixed eyes, like many insects and some birds, retinal image stabilization is achieved by compensatory head or body movements. Primates and many invertebrates with an efficient oculo-motor apparatus rely mostly on compensatory eye movements. As a matter of fact, the “hardware” triggering compensatory motor responses is common to many biological species. A wide range of mechano-neural transducers, functionally equivalent to rotation or translation-sensitive mechanisms, are found in many species [48]. One can speculate about the advantages of these particular motion sensing “transducers”, but nevertheless it remains that such a particular design solution has been naturally selected to deal with the image stabilization problem. In primates, the mechanism controlling the direction of gaze on the basis of inertial information is called *vestibulo-ocular reflex* (VOR). It is subdivided into *angular VOR* (AVOR)—generating oculo-motor responses to *angular* head motion—and *translational VOR* (TVOR)—generating responses to *linear* head motion [30,36]. In the case of the AVOR, the sensing is performed by three ring-shaped sensors (called semi-circular canals), measuring angular velocities along three perpendicular directions. In the case of the LVOR, the sensing is performed by the otoliths organs, which sense linear movements in horizontal and vertical directions and orientation of the head with respect to gravity [19]. The vestibular reflexes are known to operate in open-loop, are very rapid and work best for high frequency movements of the head [5,20,48]. On the other hand, the visual reflexes, like the opto-kinetic reflex (OKR), operate in closed-loop, are slower and respond better for lower frequencies of head movements [2,24].

The use of inertial information in artificial systems was already proposed in the past by several authors (see [4,7,23,31,45]). On the other hand, only few robotics implementations have appeared exploiting inertial cues for line-of-sight stabilization (e.g. [1]). The problem of image stabilization can be approached in different ways according to the structure of the system itself and to the nature of the environment. In recent work, Terzopoulos and Rabie [44] proved that a “*pure visual*” approach is successful within their Animat world, a virtual world where autonomous

virtual robots can perceive and interact in a physics-based environment. For the purpose of this paper though, the only interesting situation is characterized by real cameras moving in a dynamic environment. In this condition *stabilization* means being able to track a point in the environment and not simply being able to maintain a stable orientation of the camera’s optical axis. This goal can be achieved, in principle, through a process acting only on image features (see e.g. [11,28,38,47]). On the other hand, the goal of this work is to demonstrate that a more efficient solution is based on inertial measures providing angular and linear egomotion information independent of visual egomotion information derived, for example, through the optical flow field [15,16,22]. The inertial data, even if not sufficient to solve the stabilization problem, reduces the computation load of visual processing.

The paper is organized as follows. Section 2 describes the inertial and visual sensory systems of the LIRA head. Section 3 outlines some basic kinematics of gaze for translational and rotational movements. Section 4 proposes two control schemes developed in the framework of stabilization control: the first, with the intent to realize a monocular image stabilization, the second oriented to binocular systems. They both rely on the integrated use of inertial and visual information. In Section 5 the experimental results and the stabilization performance of the LIRA robot are presented. The stimuli used to characterize the system’s response are externally imposed rotational and translational movements. Finally, the advantages of using a visuo-inertial stabilization mechanism in the context of fast coordinated eye-head movements are also addressed.

2. Inertial and space-variant sensing

The inertial sensory system for the robot head has been entirely designed and realized at LIRA Lab. This sensory module is designed according to the principles of the corresponding biological sensor, the vestibular system. The device is able to measure two angular velocities and two linear accelerations along two perpendicular directions. It is composed of modular, independent, sensing elements packaged together with the electronics for filtering and signal conditioning. The sensing modules are positioned as indicated in

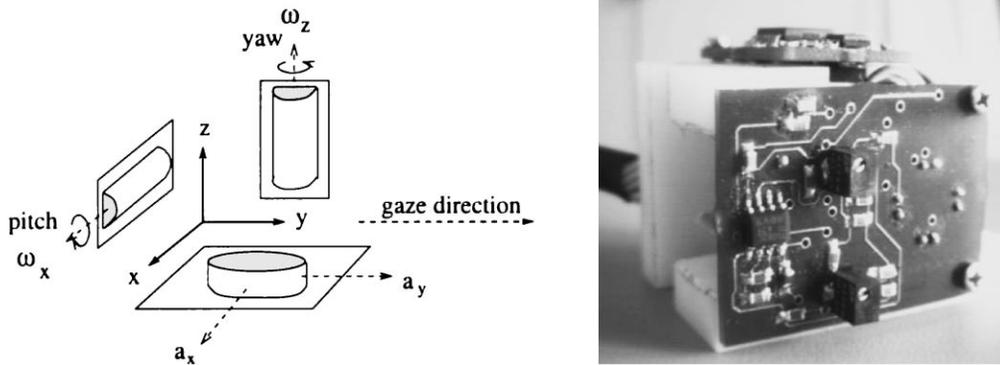


Fig. 1. The artificial vestibular system. *Left*: the spatial arrangement of the rotational and translational sensing elements. *Right*: the prototype artificial vestibular system. Overall dimensions are kept relatively small ($5 \times 4 \times 4$ cm) by implementing the design of the individual modules using surface mount technology.



Fig. 2. The space-variant image geometry. *Left*: log-polar image of 128×64 pixels. *Right*: the corresponding reconstructed Cartesian image.

Fig. 1 (left); on the right side, the prototyped device is shown. The overall dimensions are $5 \text{ cm} \times 4 \text{ cm} \times 4 \text{ cm}$. The rotational part is composed of piezoelectric sensors (also known as vibrating gyros) produced by Murata. These transducers can measure rotational velocity around their longitudinal axis: the range of measurement extends from -90 to $+90^\circ$ per second, the bandwidth from DC to 50 Hz (90° phase delay). The output of the sensor is band-pass filtered (0.03–5.0 Hz) and amplified by custom electronics. The linear part, on the other hand, is composed of sensing elements exploiting a heat transfer design principle and are produced by Pewarton. The output generated is propor-

tional to linear acceleration. The range of measurement extends from -1 to $+1$ g. In the linear part too, custom electronics is used to low-pass filter and amplify the response of the sensing elements [32].

As for the visual sensing, the LIRA's robot head acquires and processes images in a space-variant format (see [34] for a detailed description of the log-polar, space-variant geometry). Instead of a uniform resolution across the whole visual field, the robot eyes observe the environment through *foveal* image area of high resolution and a peripheral image area of low resolution. Fig. 2 (left) shows a space-variant, log-polar image as acquired by the head's vision system. The

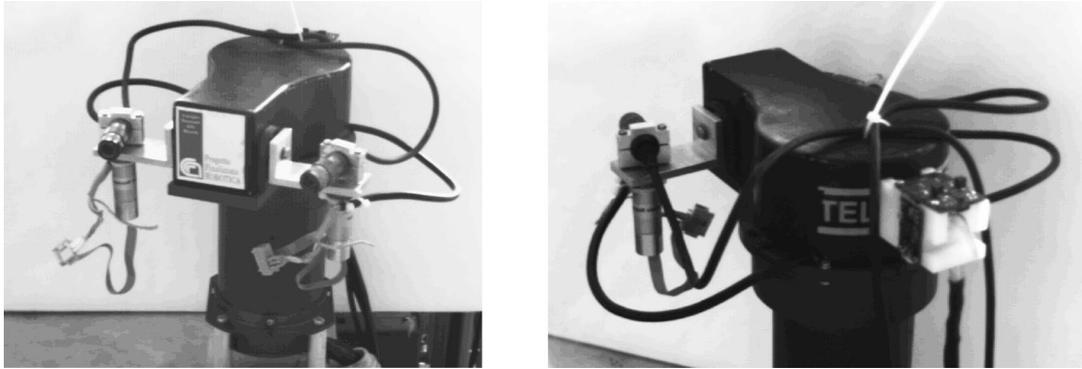


Fig. 3. The head mount. *Left*: front view of the robot head shows the offset position of the cameras with respect to the neck rotational axis (head geometrical parameters, $b = 136$ mm, $a = 95$ mm). *Right*: the artificial “vestibular” system is rigidly fixed on the back of the head mount.

image (128×64 pixels) is magnified to better illustrate the resolution variability from foveal to peripheral areas. Fig. 2 (right) displays the corresponding Cartesian reconstructed image. These two biologically inspired sensory systems are integrated within a binocular architecture (see Fig. 3). More specifically, the active vision hardware used in these experiments is composed of:

- a four d.o.f. mount, independent vergence, common elevation and common pan,
- two color cameras simulating a space-variant (log-polar) geometry,
- an artificial vestibular system,
- a double Pentium-Pro system architecture, a frame grabber (IntegralTech Flash Point 128) and a motion control board (Motion Engineering).

3. Head movement versus gaze kinematics

The goal of this section is to make explicit the relative roles of inertial and visual information in different situations. It should help to highlight the advantages of visuo-inertial integration and the asymmetries of eye-control commands. The description we develop here includes the eye-head geometric parameters and the eye-head kinematic dependence on fixation distance for rotations and translations. This formalism should not be taken to suggest that biological systems explicitly compute direct or inverse kinematics.

3.1. Rotational movements

We consider here the compensatory eye movements required to maintain stable fixation of a target at distance d , when the head rotates around a vertical, off-centered axis. Fig. 4 shows the geometry of a binocular system for this case and indicates the most relevant geometrical parameters: the inter-ocular distance (or baseline) b , the perpendicular distance a between the rotational axis of the head and the baseline, and the viewing distance d , measured from the head rotational axis and the gaze point. The analytical relation among these parameters can be derived by considering the kinematics of this model, and imposing the constraint that the eye E maintains gaze at point P when the head rotates (see Appendix A.1 for details). Simple vectorial rules and differentiation with respect to time leads to the following expression of angular velocity ω_1 :

$$\omega_1 = \left[1 + \frac{dZ_1 - (a^2 + \frac{1}{4}b^2)}{d^2 - 2dZ_1 + (a^2 + \frac{1}{4}b^2)} \right] \omega_h, \quad (1)$$

where $Z_1 = (a \cos \theta_h + \frac{1}{2}b \sin \theta_h)$ represents the Z -coordinate of the left eye. Eq. (1) determines the relationship between eye velocity, ω_1 , and (i) the geometrical parameters of the eye-head system (i.e. b and a) and (ii) the distance d of the fixation point P for any given head velocity ω_h . In a robot vision system different choices of the a and b parameters (i.e.

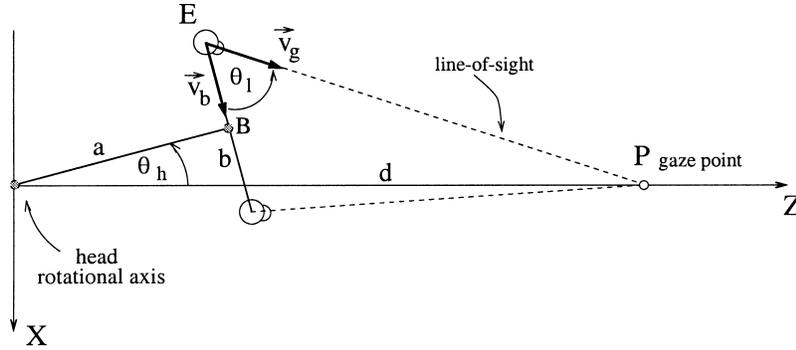


Fig. 4. Geometry of the eye-head system showing the parameters (a, b, d) relevant to inertial and visual measures.

geometric configurations and dimensions) determine different shapes of the eye-head velocity relationship. Eq. (1) also makes explicit the inverse dependence upon distance [33]. Interestingly enough, dependence upon distance is clearly evident in primate's RVOR responses [6,13,17,40,43,46]. These findings suggest that fixation distance might play an important role in simplifying/synthesizing efficient oculo-motor responses to rotational movements especially in the close range domain. In this regard, note that the eye velocity ω_l required to maintain fixation on near objects can be as much as twice the value of ω_h for fixation distances in the range 25–200 cm; in other words, the optimal amount of ocular compensation needed to achieve a stable gaze can change rapidly with distance of fixation. One point worth stressing here is the fact that the range over which fixation distance influences the optimal control of the eye compensatory gain (at least from the kinematics point of view) may not be very relevant for locomotion (e.g. a robot walking/navigating and fixating at a long range). On the other hand, this range overlaps entirely with manipulation workspace, and in this respect, might justify appropriate control circuits (e.g. a robot manipulating objects).

3.2. Translational movements

To understand the requirements of stable fixation in the case of translational movements, we refer to the same formalism used for rotations. Fig. 5 shows the geometric representation of this task for a binocular system which is translating with instantaneous veloc-

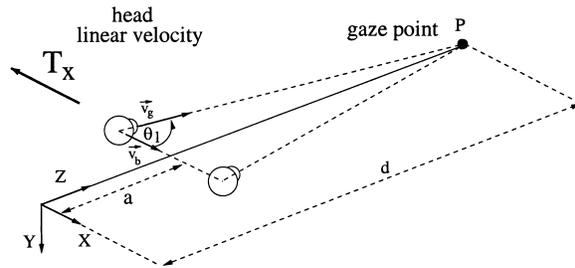


Fig. 5. Geometry of ocular responses to lateral translations.

ity T_x along the x -axis. The vectorial formalism applied in this case gives

$$\omega_l = \left[\frac{(d - a)}{(\frac{1}{2}b - x)^2 + (d - a)^2} \right] T_x. \quad (2)$$

Fig. 6 represents graphically the expression of the gain required for perfect compensation (i.e. Eq. (2)). The curve shows clearly that the ocular response required to maintain fixation on near objects (i.e. 20–150 cm) can be quite large, but decreases inversely with distance. An object at infinite distance does not require, in principle, any ocular compensation irrespective of how fast translation is occurring. Moreover, when the eye is fixating an object in the range 50–200 cm, the required gain (i.e. required eye velocity per unit linear translation) changes with distance dramatically: from a value of 0.5 at a distance of 150 cm, the gain rises to about 2.0 at 50 cm. In terms of ocular velocity, this means that changing fixation from a point at 150 cm distance to another at 50 cm, while translating at velocity T_x , requires a fourfold increase in the eye

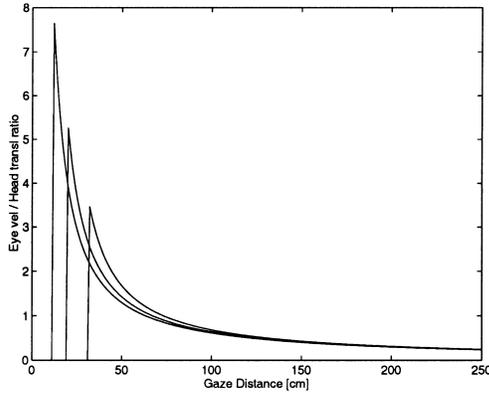


Fig. 6. Theoretical eye-head velocity ratios plotted for translations of the head. The velocity ratio as a function of fixation distance d for different values of a . The baseline b is fixed (6 cm) and eye-to-neck distance a increases (from 6 to 16 cm).

velocity if fixation is to be maintained. Once again, it is interesting to remark that the gain of the primate TVOR and RVOR varies inversely with viewing distance, although compensation is often less than complete. The modulation of the TVOR gain is subject to instantaneous changes in fixation distance and it can be modeled using a linear relationship with binocular vergence angle, though vergence is not the only parameter used for range-finding [37,39].

3.3. Eccentric gaze

When looking at an eccentric target in near space, the compensatory eye movements required to maintain binocular alignment during head rotations are different for the two eyes [17]. From a kinematics point of view, the origin of this asymmetry is clear if one compares the analytical expressions of the two angular velocities (i.e. left vs. right eye). The optimal compensatory velocities are given by Eq. (1) for the left eye and by the following for the right eye:

$$\omega_r = \left[1 + \frac{dZ_r - (a^2 + \frac{1}{4}b^2)}{d^2 - 2dZ_r + (a^2 + \frac{1}{4}b^2)} \right] \omega_h, \quad (3)$$

where $Z_r = (a \cos \theta_h - \frac{1}{2}b \sin \theta_h)$ represents the Z-coordinate of it. The expressions are almost identical except for the sign of a few terms encoding the opposite position of the eyes. The representation of the two gains in polar coordinates (see Fig. 7—gains are plotted with respect to head angular position, θ_h , and for a given distance of fixation) reflects the asymmetric requirement of the optimal response. In the case of an object at 30 cm distance (Fig. 7, left), an angular deviation from the frontal direction of 30° introduces a relative gain difference between the two eyes of about 0.2; for example, with a head velocity of 200° per second, deviating 30° from the frontal

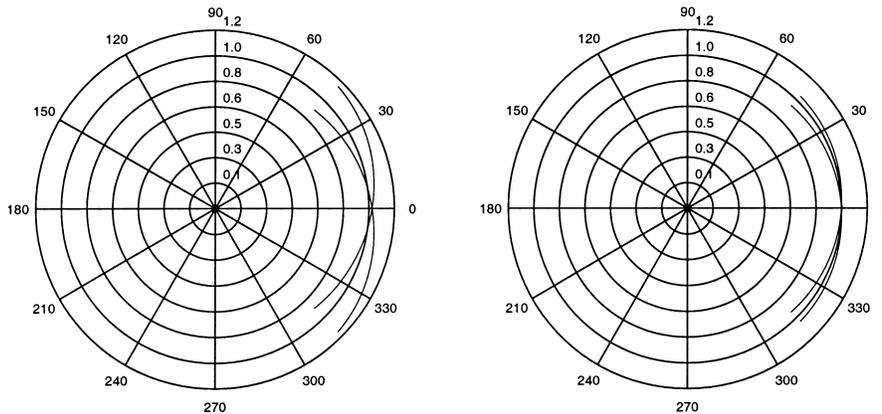


Fig. 7. Polar representation of the inertial gains (G_{VOR_l} and G_{VOR_r}) for the left and the right eye. The two curves in each plot describe the binocular gains required to perfectly stabilize the left and right image during a rotational movement. The angular displacement from the zero reference codes the orientation of the head, while fixating frontally at a given distance. *Left*: distance of fixation is 30 cm. *Right*: distance of fixation is increased to 70 cm.

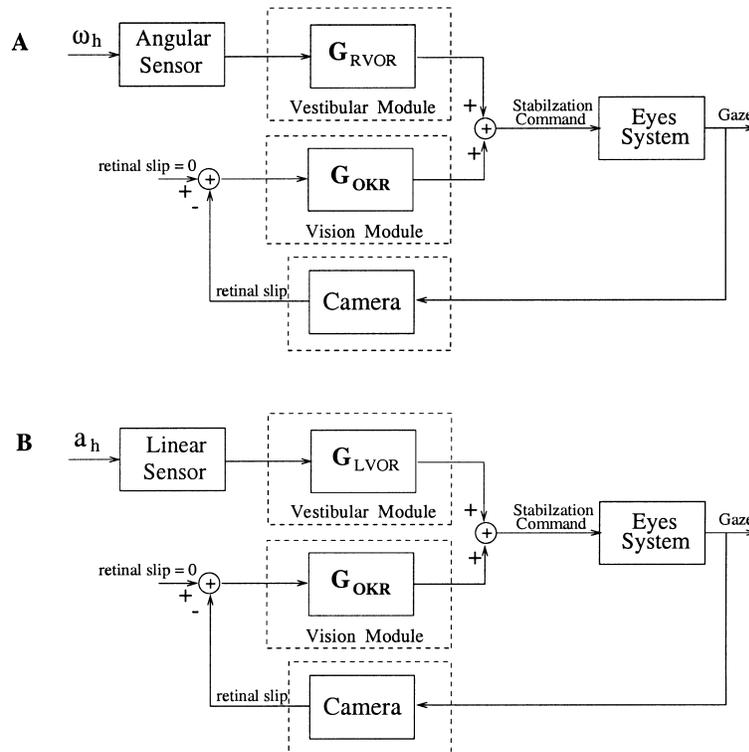


Fig. 8. Block diagrams of visuo-inertial mechanisms to stabilize robot's gaze. (A) The rotational module selectively measures rotational movements of the head (ω_h). (B) The linear accelerometer senses translational movements of the head (a_h). In both cases, the inertial information is processed open-loop and acts in parallel way with the vision module.

direction gives a relative angular differential velocity of 40° per second. Thus, the angular velocities of the two eyes can be rather different in the near space. Although in humans, there is clear evidence that during compensatory eye movements, binocular alignment is not strictly maintained [10], this constraint might be more important for a robot vision system, especially if the system uses binocularly derived cues to control camera movements [9].

4. Visuo-inertial integration

The issue of synthesizing simple and effective control strategies that integrate visual and inertial information appears challenging. In an anthropomorphic binocular system, the rotational axis of the eye and of the head do not coincide. Therefore, a rotational movement of the head causes both a rotation and a trans-

lation of the eyes. If the fixation point is at infinity, the inertial information alone will in principle be sufficient to stabilize gaze perfectly. However, this situation might represent a minority of cases. In general, as we have outlined in Section 3, to generate an adequate compensatory response, the information about viewing distance is required. We suggest later in this section the use of a control strategy exploiting range information to tune the stabilization performance. Moreover, compensatory eye movements for translations also depend on other contextual factors such as the direction of gaze with respect to the direction of heading [18].

The structure of the basic visuo-inertial stabilization mechanism implemented on the LIRA head is sketched in Fig. 8. In the block diagrams, the visual and the inertial information are simply added together (in the biological literature this early modeling was named the "linear summation hypothesis"). The compensatory movement generated by the inertial infor-

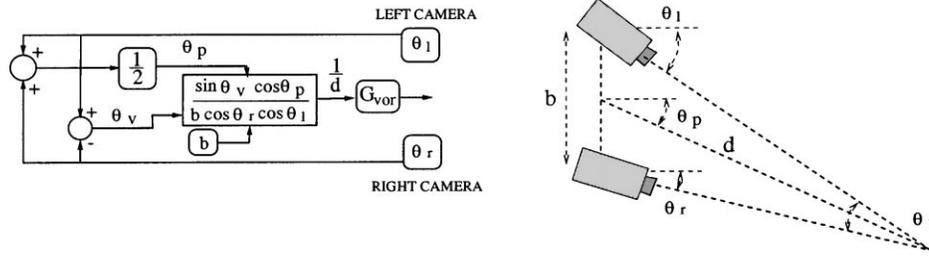


Fig. 9. Block diagram describing the computational scheme used to derive distance information.

mation limits the amplitude of image motion to a range of values measurable by visual algorithms. Optic flow is therefore used to measure the residual error (called residual optic flow — ROF) present in each image after inertial compensation. In particular, with reference to Fig. 8(A), the stabilization command generated is

$$SC(i+1) = G_{RVOR} \cdot v_h(i) + G_{OKR} * ROF(i),$$

whereas in the linear case (Fig. 8(B)) the stabilization command is given by

$$SC(i+1) = G_{LVOR} \cdot v_h(i) + G_{OKR} * ROF(i),$$

where $v_h(i)$ term represents an estimate of translational velocity as obtained from numeric integration of a finite number of $a_h(i)$ samples. Even if simply realized, this stabilization scheme, and the experimental result obtained from it, draw the attention to the following important issues. First, the computational resources in an autonomous system are not unlimited. Second, the visual processing is computationally demanding: the amount of image data is usually much larger than many other sensory data; further, vision algorithms require, on the average, several pre-processing steps. In addition to that, whatever be the algorithm adopted for image slip estimation, a limited range of image velocities can be reliably measured (we have estimated this limit for our system and results are summarized in Appendix A.2). As a consequence, in a “real-world” environment, an approach to stabilization willing to consider visual information alone will have a relatively limited performance. In contrast, the control scheme described here enables the robot head, for example, to maintain a stable gaze at 1 m distance in response to transient rotations of the head at a speed of 10° per second, necessitating the eye/camera to rotate at 12° per second for full compensation; the

inertial mechanism ($G_{VOR} = 1$) accounting for 10° per second of this compensation, leaving the visual compensation mechanism to deal with a mere 2° per second residual flow. The advantage obtained is therefore clear: the use of inertial information enables more efficient use of the computational resources and it extends the range of motions or external disturbances the system can effectively deal with, without increasing system complexity. In biological systems the same sort of trade-off is present: vestibular reflexes have high-pass dynamics and are efficiently complemented by visual reflexes with low-pass dynamics [29,42]. Just to give a numerical example, experimental data on primate vestibulo-ocular responses indicate that they are ultra-rapid (RVOR latency is <10 ms [41], and TVOR latency is <20 ms [8]), while the complementary visual responses are considerably slower (>50 ms in monkeys [25] and >80 ms in humans [14]).

In order to work optimally also in a context where the fixation point is changed in depth, a tuning scheme which adapts the compensatory gains to changes in distance is proposed. The scheme exploits range information which is derived on the basis of dynamic control of vergence [9]. This type of control enables the system to keep both eyes/cameras pointing at the object of interest and to dynamically change their orientation in order to maintain the fixated object centered in the foveas. With reference to Fig. 9, the distance d of the object being fixated can be derived from the vergence (θ_v) and version (θ_p) angles using the following equation:

$$d = b \frac{\cos(\theta_r) \cos(\theta_l)}{\sin(\theta_v) \cos(\theta_p)}, \quad (4)$$

where $\theta_p = \frac{1}{2}(\theta_l + \theta_r)$ and $\theta_v = (\theta_l - \theta_r)$. Distance information is in turn used to adaptively change the

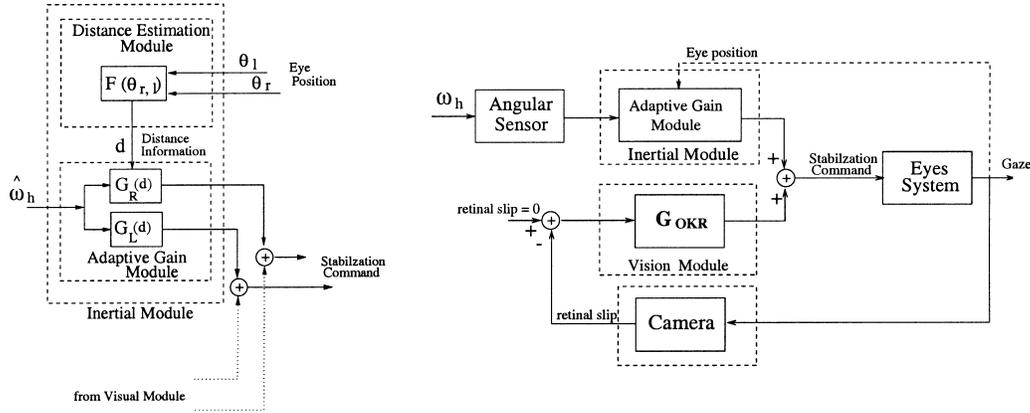


Fig. 10. Adaptive tuning of compensatory gains. *Left*: the compensatory gain of the left eye, G_L , and the corresponding gain for the right eye, G_R , are tuned according to distance of fixation. The block containing the function $F(\cdot)$ combines eye angular positions to derive an estimate of the fixation distance d according to expression (4). *Right*: the adaptive gain module replaces the constant gain module of the previous scheme.

compensatory gains of both eyes. The diagram in Fig. 10 describes the information flow which implements this scheme. The tuning functions we propose to use at this stage are those obtained from the kinematic analysis of Section 3.3. A different approach would be that of learning these tuning functions.

5. Experiments

Four different experiments are presented. The first and second ones describe the stabilization performance in the monocular case: rotational movements around a vertical axis and linear translations along a fronto-parallel plane are addressed. The translational case will be described more in detail. In fact, the stabilization performance for the rotational case has been extensively addressed in a previous work [33]. For both types of movements, the results show that slip on the image plane can be reduced and almost canceled by appropriate, constant tuning of the inertial (G_{VOR}) and visual (G_{vis}) gains. The third experiment presents an adaptive tuning scheme for the eye compensatory response for the rotational case. The robot system estimates the distance of the fixation point and uses this parameter to tune the gains of the counter-rotational response. The last experiment shows to what extent the use of inertial information can improve image stability during coordinated eye–head movements. In all

the experiments, stabilization performance is evaluated quantitatively by using one (in some case two) of the following visual measurements: correlation measurement between consecutive frames (also called image stabilization index — ISI), retinal error and residual optic flow (ROF). A detailed description of the technique used to estimate the ROF (from log-polar images) is given elsewhere [9]. On the other hand, the ISI index is defined as follows: $ISI_i = 1 - NC_i$,

$$NC_i = \frac{\sum_{\eta, \xi} (I_i(\eta, \xi) - \mu_i) \cdot (I_{i-1}(\eta, \xi) - \mu_{i-1})}{\sqrt{\sum_{\eta, \xi} (I_i(\eta, \xi) - \mu_i)^2 \cdot \sum_{\eta, \xi} (I_{i-1}(\eta, \xi) - \mu_{i-1})^2}} \quad (5)$$

The correlation measures the degree of similarity between two subsequent log-polar images (I_i and I_{i-1}). The symbols μ_i and μ_{i-1} indicate the corresponding image mean values. Better stabilization performance is mapped to lower values of ISI. During each experiment the following data are stored to enable an off-line evaluation of the overall performance: (i) position angle of the inertial stimulus; (ii) output of the inertial sensors; (iii) the translational component of the optic flow; (iv) values of the image stabilization index (ISI); (v) position angle of the controlled camera; (vi) retinal target position with respect to the center of the image (this index is used in

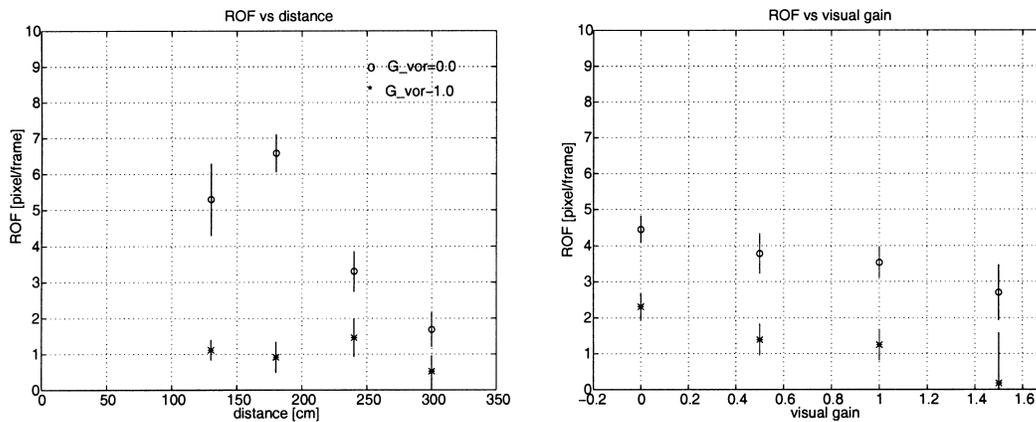


Fig. 11. Residual optic flow (ROF) on the image plane. *Left:* ROF for different distances when two different G_{VOR} gains are used. *Right:* ROF for different values of visual gain G_{vis} (G_{VOR} is constant: the “○” symbol indicates measurements corresponding to $G_{\text{VOR}} = 0.7$, while the “*” symbol indicates $G_{\text{VOR}} = 1$).

particular in the binocular stabilization experiment). Position and inertial information are sampled at a rate of 33 Hz (i.e. 30 ms interval). All the visual processing is performed at a sampling interval of 80 ms (i.e. 12.5 Hz). As regards the visual parameters measuring stabilization performance, the values reported in all the graphic presentations are obtained as statistics over several repetitive trials; for this reason data are often indicated along with the corresponding standard deviation.

5.1. Monocular stabilization: Head rotations

The results of monocular image stabilization during head rotation around the vertical axis are here only summarized [33]. In this experiment the rotational axis of the inertial stimulus coincides with the common pan axis of the camera mount. The stabilization performance is evaluated in terms of the ROF measured at the center of the image plane. The parameters of the rotational stimulus (amplitude and frequency) were chosen to induce on the image plane a motion field close to the saturation threshold of the visual processing system. In our case this threshold is of the order of 4–5 pixels. Above this value, the relation between the image velocity and the optic flow is not linear (see Appendix A.2). It is worth stressing that even if the dynamics of the stimulus is not a “high dynamics” (i.e.

10° per second amplitude, 0.2 Hz frequency), in this case the pure visual approach could not be performed correctly (i.e. the optic flow processing is out of range). This experiment is performed in real-time with a control period of 80 ms.¹ Two sets of ROF measurements are presented. The first, plotted in Fig. 11 (left) shows the decrease of the ROF with and without inertial stabilization as function of fixation distance. Note that in the case $G_{\text{VOR}} = 1$, the ROF measured is always maintained within the range of correct measurements. It is in this case that ROF can be used to synthesize an additional motor signal to cancel out any remaining image slip. On the other hand, when the inertially generated response is purposefully removed ($G_{\text{VOR}} = 0$), the values of the ROF obtained for a distance of 130 cm start to indicate saturation effects (i.e. the ROF drops to a lower value and the corresponding standard deviation indicates a larger variability of the measurement) generating unreliable information for the feedback loop. The second set of measurements (Fig. 11, right) show that when visual and inertial information are integrated, the performance of stabilization further improves. The two series of ROF measurements show a constant trend toward lower values when the visual gain G_{vis} is introduced.

¹The architecture used in this experiment did not support multi-tasking; this issue constrained the sampling interval of visual and inertial information to a common value.

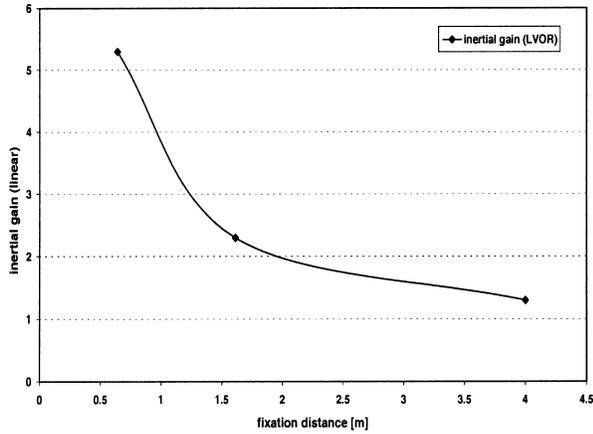


Fig. 12. Linear inertial gain as a function of distance. Gain values are determined experimentally by minimizing the ISI visual measurement.

5.2. Monocular stabilization: Head translation

This experiment describes the image stabilization performance of the LIRA robot in response to horizontal movements in the fronto-parallel plane. These types of movements elicit the response of the linear part of the artificial “vestibular” system. Compensatory eye movements for fixation in the fronto-parallel perpendicular plane are addressed. The robot head is mounted on a support which is fixed to a computer controlled, translating table. This experimental arrangement enables to generate repetitive left-and-right movements with a desired acceleration profile. For the purpose of the current measurements, the cameras of the head are initially oriented to fixate a background plane; distance to the background, which is an experimental parameter, is changed throughout trials. The motion profile generated by the motorized table is trapezoidal, with an acceleration of 90 cm/s^2 and a maximum velocity of 30 cm/s . Three different fixation distances are considered: 64, 161 and 400 cm. For each distance an “optimal” tuning of the “linear” gain was obtained experimentally by minimizing the ISI index. The values determined for the gains are in the corresponding order, $G_{\text{VOR}} = 5.3$, 2.3 and 1.3 . In Fig. 12 we have plotted such gain values as a function of the distance parameter (i.e. distance to the background plane). Interestingly enough, the curve interpolating the experimental gains shows the same inverse dependence on

distance as the “kinematic” curve derived in Section 3.2. For each experimental trial, performance in stabilization is evaluated by calculating the residual optic flow (ROF) corresponding to the maximum translational velocity. Fig. 13 shows three ROF curves corresponding to the three different distances mentioned above. The left-most value in each curve represents the optic flow measured when no compensatory response is generated. Clearly enough, the higher the fixation distance (bottom value, $d = 4.00$), the lower the amplitude of the measured flow. When compensatory eye movements are generated (with the “optimal” linear gains), the ROF becomes smaller. Starting from each second value, the samples represent the ROF measured using the “optimal” inertial gain G_{VOR} and an increasing value of the visual gain G_{vis} (range of the visual gain used extends from 0.01 to 0.10). Although in two cases, namely $d = 400 \text{ cm}$ and $d = 161 \text{ cm}$, stabilization performance improves with the appropriate tuning of the visuo-inertial response, the top curve ($d = 64 \text{ cm}$) shows clearly that for small distances image motion still remains high (in the order of 1.5 pixel/frame). At this stage, the limit for such small distances is, in our opinion, due to several factors: uncon-

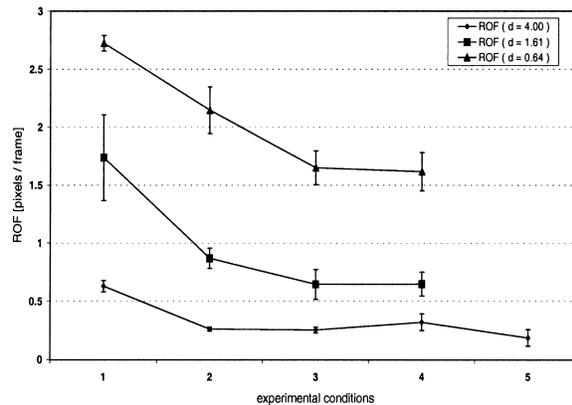


Fig. 13. Residual optic flow measured during head translations. Each curve refers to a different fixation distance (in cm) and a different G_{VOR} constant value. Top curve: $d = 64$, $G_{\text{VOR}} = 5.3$. Middle curve: $d = 161$, $G_{\text{VOR}} = 2.3$. Bottom curve: $d = 400$, $G_{\text{VOR}} = 1.3$. The experimental conditions correspond to a constant inertial gain and a different value of the visual gain for each fixation distance curve (from $d = 400$ to 64 , respectively). Conditions are: (1) $G_{\text{VOR}} = 0.0$, $G_{\text{vis}} = 0.0$ for all the curves, (2) constant G_{VOR} and $G_{\text{vis}} = 0.0$, (3) constant G_{VOR} and $G_{\text{vis}} = 0.01, 0.07, 0.05$, (4) constant G_{VOR} and $G_{\text{vis}} = 0.085, 0.1, 0.07$.

trollable delays in the sensori-motor loop, phase lags introduced by sensor filtering (to eliminate the sensor noise), but most importantly, the low sensitivity of the linear sensing device. Further work is on-going to improve the linear part of the inertial sensory system.

5.3. Binocular stabilization: Adaptive gain control

These experiments show how the robot system responds by adapting the oculo-motor compensation strategy according to distance requirements. We have shown in Section 3.1 that for rotational movements of the head a stable fixation can be obtained in the short and long range by tuning the compensatory gain inversely with distance. In Section 4, we have proposed an adaptive control scheme exploiting distance information. Here we test the performance of this control strategy for several fixation distances. To obtain reliable distance information we have programmed the robot head to control its vergence (i.e. the distance of fixation point) on the basis of color information. An algorithm performing color-segmentation runs independently on each robot's eye. Position information of a target are extracted and used to keep the object centered in both images. A more general approach to control vergence would be the scheme proposed by Capurro et al. [9], which rely on binocular fusion and does not require a priori knowledge of target characteristics. On the other hand, for the purpose of our experiment, the vergence control above mentioned is sufficient (and less computationally expensive) to derive consistent angular information relative to fixation point. When the target is foveated in each eye, the depth is calculated with Eq. (4). At the next stage, using the pure kinematic relations derived in Section 3.1, the gains of the inertially driven response can be adapted "on the fly" to the necessary requirements. A stimulus of constant amplitude and frequency (17° per second and 0.3 Hz) is used to impose sinusoidal rotational movements to the head. Stabilization performance is evaluated by means of the retinal error (in pixels) and of the ROF at a sampling interval of 80 ms. Figs. 14 and 15 show, respectively, the retinal error and the ROF for the "non-compensated" and the "compensated" case. The top curves of Fig. 14 describe the error without compensatory response. Note that when fixation distance is 50 cm, the error

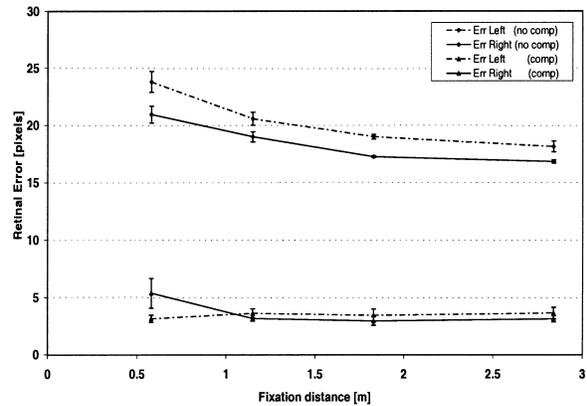


Fig. 14. Retinal error measured during head rotational movements. Upper curves measure error in the "non-compensated" case; lower curves represent error in the "adaptive compensation" case.

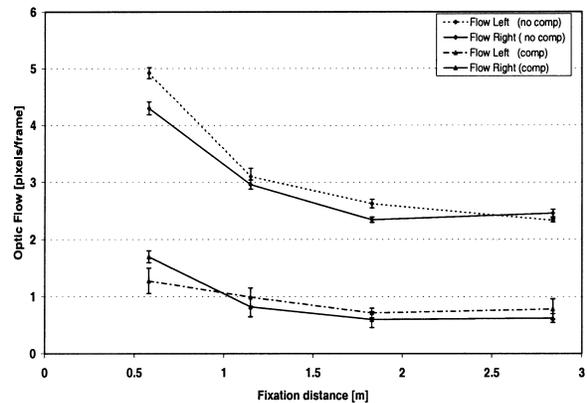


Fig. 15. Optic flow measured during head rotational movements. Upper curves: "non-compensated" case. Lower curves: "adaptive-compensation" case.

is approximately 24 pixels; for a fixation distance of 280 cm, it reduces to 18 pixels. On the other hand, in the compensated case (same figure, bottom curves) the error remains almost constant when fixation distance changes from 280 to 50 cm. These results prove that the adaptive tuning of the inertial loop is successful. The values of the adaptive gains derived autonomously by the system on the basis of vergence control are shown in Fig. 16. The second index of performance, the ROF, is shown in Fig. 15. Note that flow amplitude for fixation distance of 50 cm reduces from 4.5 pixels/frame to about 1.5 pixels/frame in the "compensated case". Furthermore, in the "compen-

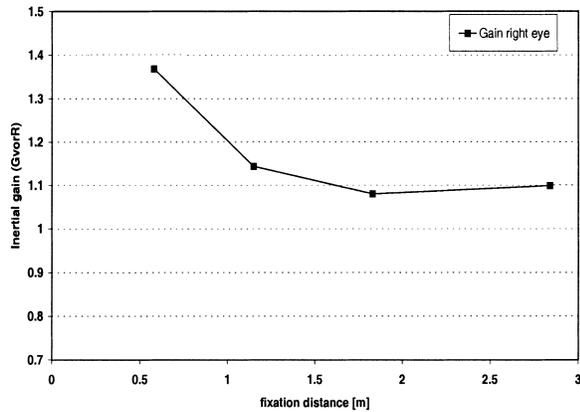


Fig. 16. Adaptive gain G_{VORr} computed autonomously by the system on the basis of control of vergence.

sated case” the variability of ROF within the whole fixation range (i.e. 50–280 cm) is smaller than in the non-compensated case. A sequence of frames acquired in the adaptive (left) and non-adaptive (right) stabilized conditions during this experiment is shown in Fig. 17.

5.4. Gaze stabilization during eye–head movements

This section describes the role of inertial information during gaze redirections. In particular, it focuses on fast, coordinated robot eye–head movements generated to foveate targets appearing at eccentric positions. The results show that when the inertial information is used the robot behavior is characterized by a faster fixation and a smaller overshoot of the target.

In a series of repeated trials a target is presented to the robot in a fronto-parallel plane at a distance of 100 cm. Three eccentric positions are chosen, respectively, at 5, 10 and 15° eccentricity. These eccentricity values represent a reasonable sampling of the robot visual field given the focal length of the cameras (7.5 mm). The movement of the eyes is programmed in retino topic coordinates by transforming image coordinates in motor commands through two look-up tables (LUTs), one for each eye. The look-up tables are filled before the experiment by using a calibration procedure. At the same time a (one-shot) head command is generated taking into account the programmed eye positions at the end of the saccade. Also the head commands are stored in a look-up table. The initial

saccadic movement of the eye is open-loop and lasts 4 cycles (the eyes acceleration is set to $50^\circ/s^2$, saccade is performed in $4 * 30 = 120$ ms interval). When the saccadic movement is over, the control of the eye movements is switched to closed-loop. Eye movement is controlled on the basis of color information. The head (i.e. the common pan axis of the cameras) follows the eyes to re-establish a symmetric eye/head configuration using a PID controller. The error signal driving the common pan axis loop codes the relative orientation of the eyes with respect to the head. The head maximum acceleration is set to $10^\circ/s^2$.

The stabilization performance is measured quantitatively in two cases: (a) the inertial information is used by the robot to generate compensatory eye movements, and (b) the inertial information is not used. The performance measurements is obtained in terms of: (1) the target overshoot, i.e. the transients target position error after the saccadic part of the eye control is completed, and (2) the time interval required for ISI to fall below the given threshold. In particular, overshoot is computed as the difference between the minimal retinal error and the maximal retinal error after the saccadic movement. The threshold for the ISI index was experimentally established to be 0.3. In fact, when ISI falls below this value processing of image feature and extraction of dynamic parameters lead to accurate and robust measurement (see optic flow, Appendix A.2, range and accuracy measurement).

Fig. 18 shows the dynamic trajectories of eye–head coordinated movement for a saccade performed at 15° angular eccentricity. On the left side, the measurements are obtained when inertial information is used. On the right side, the same measurements are derived when inertial information is not considered. From top to bottom, the trajectories represent: the left eye movement (θ_e), the head movement (θ_h), the angular velocity measured by the inertial system, the corresponding gaze angle ($g = \theta_h + \theta_e$). Fig. 19 shows the target position and the ISI index measured during the same experiment. The comparison of the data in Fig. 19 shows that: (i) in the “compensated” case (left), the overshoot of the target is smaller than 12 pixels (see the difference between the marks “*” and “+” around the time unit 15); in the “non-compensated” case the overshoot becomes as large as 25 pixels; (ii) the ISI index becomes smaller than the required threshold earlier in time (at about time unit 17 versus 27 in the

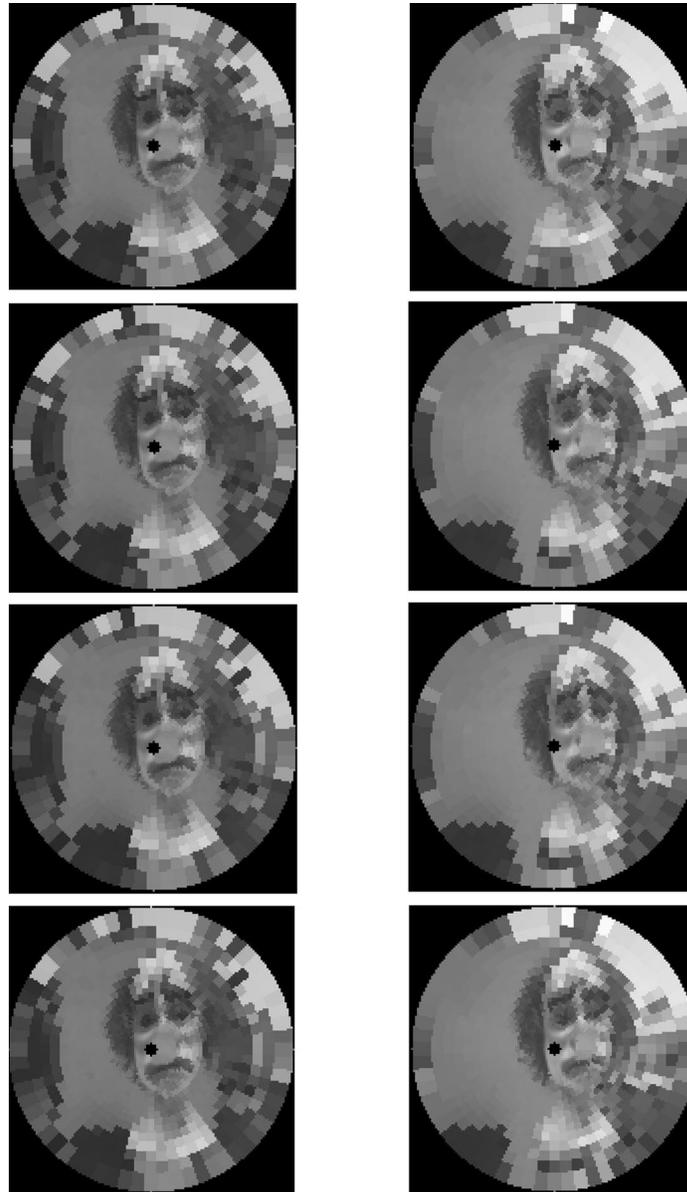


Fig. 17. Left image of stereo pairs acquired during the binocular stabilization experiment. The nose of the clown represents the target to be tracked by color processing. *Left*: images of the adaptive compensated camera. *Right*: images of the non-adaptive compensated camera. In the latter case, note the relative larger error in tracking the target feature. The space-variant images processed by the robot are shown here remapped in Cartesian coordinates.

“non-compensated” case). These measurements have been performed several times ($N = 6$) and for different saccade amplitudes, namely 5, 10, 15° eccentricity. Data have been averaged and standard deviation computed. Fig. 20 summarize the results. The overshoot

for the “compensated” case is always smaller than the corresponding “non-compensated” case (see Fig. 20, left). At the same time, the stabilization interval required for the ISI to fall below the robust threshold in the “compensated” case is considerably smaller. In

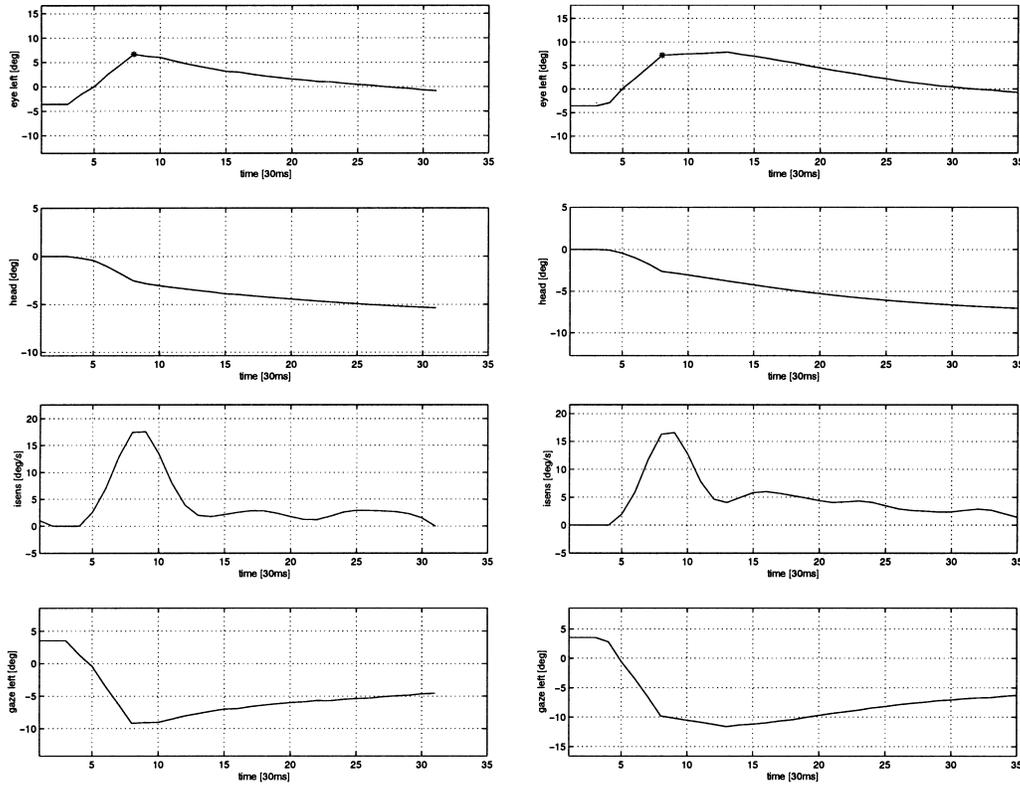


Fig. 18. Position and velocity information during two gaze redirection experiments. From top to bottom: the left eye position, the head position, the inertial sensor output and the gaze position. *Left*: the head velocity sensed by the inertial sensor is used to generate compensatory eye movements. The end of the saccadic part of the eye movement is marked with a “*” symbol. *Right*: inertial information is not used to generate compensatory eye movements.

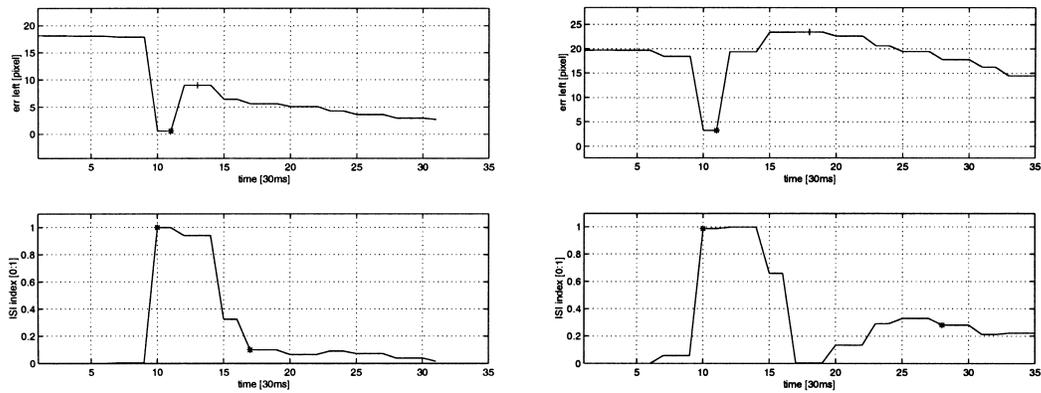


Fig. 19. Visual parameters computed during gaze redirection. *Top*: target retinal error. The overshoot is measured as the difference between the minimum (“*”) symbol and the maximum retinal error (“+”) symbol after the saccade. *Bottom*: image stabilization index (ISI). The time interval required for the ISI to fall below a stable threshold of 0.3 is delimited by two “*” symbols.

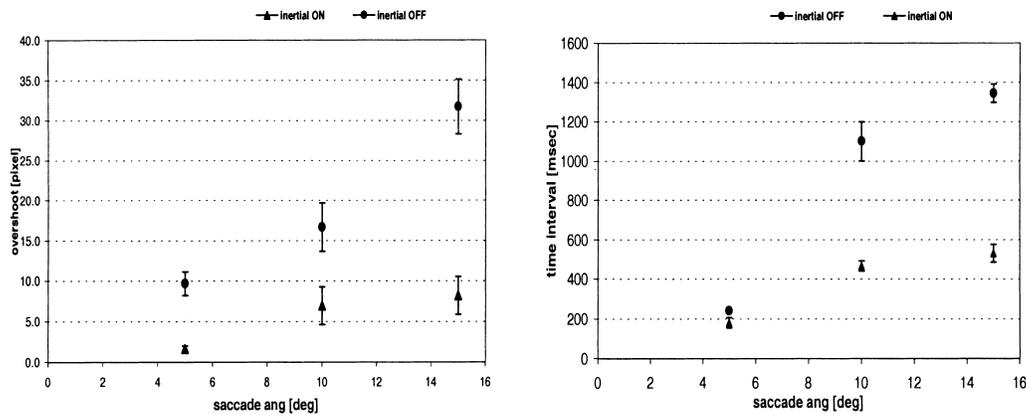


Fig. 20. Parameters measuring stabilization performance during coordinated eye–head movement. *Left*: retinal error overshoot: the inertial and the non-inertial case. *Right*: time interval required for the ISI to fall below the 0.3 threshold.

general the data shows that the use of inertial information in gaze redirection strategies leads to two important advantages: (1) an earlier stable image of the new spotted location and (2) simpler motor control strategies for gaze-line redirection involving coordinated eye–head movements.

6. Discussion

The theoretical and experimental data presented in this paper are aimed at demonstrating the following issues: (i) correct/accurate oculo-motor response for gaze stabilization (i.e. eye velocity, G_{VOR}) depends on geometrical parameters of the eye–head system and distance of fixation point; (ii) the fixation distance has important consequences for both rotational and translational head movements; (iii) the use of inertial information facilitates the visual processing; (iv) an adaptive tuning of the compensatory oculo-motor response can further reduce the visual processing requirements; (v) the use of inertial information in gaze redirection strategies increases stabilization performance in dynamic contexts and simplifies coordinated eye–head motor control.

An important point worth stressing is the difference in latency of the two “sensory” systems. Latencies depend on the amount of sensor-generated data (much larger in the case of visual sensors) to the computa-

tional complexity of algorithms used to extract useful sensory cues, and to the available hardware. Inertial data are quantitatively some orders of magnitude less than visual data. As an example, think of a robot sampling three rotational and three linear accelerometers at 1000 Hz (which is certainly much more than what the human visual system does) and acquiring and processing stereo images at 25 Hz. The amount of inertial data that needs to be processed per second is almost negligible with respect to the amount of visual data that need to be processed for reliable optical flow estimation. As a direct consequence, gaze stabilization based on visuo-inertial sensory information is more responsive: the number of sensorial sources is increased, but to some extent, the computational resources required are optimized, and as a whole, reduced. Furthermore, the inertial data do not depend on visual processing (and vice versa) and their integration does indeed add a completely new and independent data source. This is particularly true in case of external disturbances where even the proprioceptive information measuring head rotation is not relevant. In this case inertial information seems the only source of extra-retinal signals and a powerful tool to increase the overall performance of the system. As regards the image processing for retinal slip estimation, algorithms adopted for optical flow calculation do indeed have a limited range of image velocities that can be reliably measured. Consequently, pure visual stabi-

lization performance is bounded. Whatever this range is, inertial stabilization improves the performance of the system.

An additional important point has been demonstrated in the binocular stabilization experiment. Binocular vision systems can be more efficient in gaze stabilization in at least two senses: (i) If a robust control of vergence exists, they can derive consistent depth information and use it to tune the eye compensatory movements. This in turn enables the system to obtain more stable images during translation and rotations at different fixation distances. (ii) Moreover, a binocular system can exploit binocular constraints (like the horopter) to perform motion analysis limited to the visual area which is interesting to stabilize, i.e. the fixation plane (i.e. the plane in space corresponding to intersection of the optical axes, also known as zero disparity plane).

Finally, the experiment on the coordinated eye–head movements has shown the following three main points: first, fixation over-shoot produced by the head movement is considerably reduced. Second, earlier stable fixation of the new spotted location (or target) is obtained. In this regard it is worth noting that algorithms for motion estimation need at least a few stable frames to compute motion parameters consistently. Therefore, the earlier the images are stable, the sooner the “visual functions” of the robot are re-established. This, in turn, increases the robot’s reactive behavior to sudden changes occurring at the new spotted location. Third, motor control strategies for coordinated eye–head movements to eccentric targets can be simplified. It is not surprising that biological systems exploit such strategies to increase efficiency of stabilization and to simplify “neural control” of the motor plant during gaze redirection [26]. Further work will concentrate on improving the response of the linear module and on integrating it into the adaptive control scheme.

Appendix A

A.1. Angular velocity of the eye

Consider Fig. 4 where two vectors, namely \mathbf{v}_g and \mathbf{v}_b , indicate for the left eye, respectively, the

line-of-sight (i.e. the vector connecting the eye position E with the gaze point P) and the semi-baseline orientation (i.e. the vector connecting eye position E and the mid-baseline point B). The instantaneous eye angle θ_e can be expressed according to the following:

$$\theta_e = \arctan \left(\frac{\mathbf{v}_b \times \mathbf{v}_g}{\mathbf{v}_b \cdot \mathbf{v}_g} \right). \quad (\text{A.1})$$

Substituting the vector symbols with their expression in terms of a , b , d , θ_h and constraining gaze points to the Z axis, we get

$$\mathbf{v}_g = \left(\frac{1}{2}b \cos \theta_h - a \sin \theta_h, d - \left(a \cos \theta_h + \frac{1}{2}b \sin \theta_h \right) \right), \quad (\text{A.2})$$

$$\mathbf{v}_b = \left(\frac{1}{2}b \cos \theta_h, -\frac{1}{2}b \sin \theta_h \right), \quad (\text{A.3})$$

$$\theta_e = \arctan \left(\frac{2(a - d \cos \theta_h)}{b - 2d \sin \theta_h} \right). \quad (\text{A.4})$$

Differentiating Eq. (A.4) with respect to time and introducing the two auxiliary expressions $Z_l = (a \cos \theta_h + \frac{1}{2}b \sin \theta_h)$ and $Z_r = (a \cos \theta_h - \frac{1}{2}b \sin \theta_h)$ representing, respectively, the Z -coordinates of the left and right eye, we have

$$\omega_l = \left[\frac{d(d - Z_l)}{d^2 - 2dZ_l + (a^2 + \frac{1}{4}b^2)} \right] \omega_h \quad (\text{A.5})$$

for the left eye and

$$\omega_r = \left[\frac{d(d - Z_r)}{d^2 - 2dZ_r + (a^2 + \frac{1}{4}b^2)} \right] \omega_h \quad (\text{A.6})$$

for the right eye. Eq. (A.5) can also be rewritten as

$$\omega_l = \left[1 + \frac{dZ_l - (a^2 + \frac{1}{4}b^2)}{d^2 - 2dZ_l + (a^2 + \frac{1}{4}b^2)} \right] \omega_h, \quad (\text{A.7})$$

in which the inverse dependence upon distance d is made more explicit.

A.2. Optic flow: Range and accuracy measurements

For the purpose of our measurements a first-order approximation (i.e. affine model) of the optic flow is sufficient and easier to compute than local optic flow [21]. The technique applied to the space-variant image

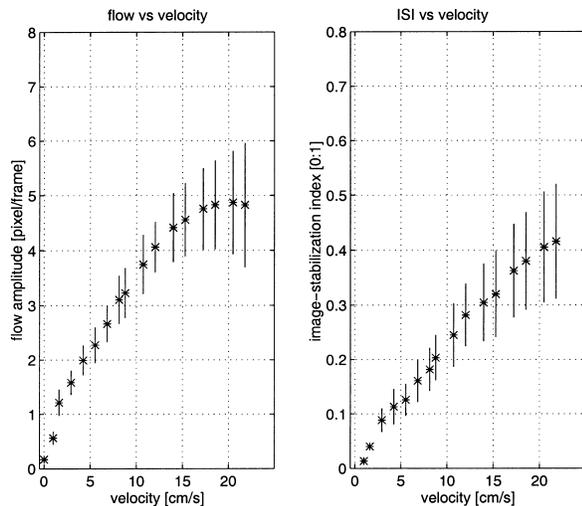


Fig. A.1. Range and accuracy of the optic flow horizontal component u_0 as measured by our algorithm: saturation effect (left), and image stabilization index (right). An ISI value above the 0.3 threshold indicates the beginning of the saturation process in the optic flow measurement.

framework is detailed in [9] and revised in [33]. Since we are going to use this motion estimation technique to evaluate the performance of the stabilization control, it is important to know quantitatively the range and the accuracy of the first-order optic flow estimates. In particular we are interested in the u_0 horizontal component of optic flow. We have derived this information experimentally by repetitive measurements: a camera fixed on top of a controllable slide is repetitively translated backward and forward at constant velocity, while the u_0 estimates are collected. The camera is moving in front of a textured flat surface, 245 cm distant, along a fronto-parallel trajectory. At the end of each back-and-forth motion sequence, the velocity of the slide is increased. The range of translation velocities is between 1 cm/s and more than 20 cm/s. Fig. A.1 shows the u_0 component (left-side) and the image stabilization index (ISI) (right-side) as a function of translation velocity. The mean value and the standard deviation are plotted for each experiment. The data show a good behavior of the u_0 estimate for translation velocities up to 17 cm/s. For higher translation velocities (more than 20 cm/s) the u_0 estimate saturates. These recordings, once again, emphasize in a quantitative way that for a given amount of computational

resources, visual information alone can be used reliably only within the “good range” of measurements, i.e. up to the point where a “saturation effect” starts corrupting the accuracy and robustness of the visual measurement process. In Section 6, this point is further elaborated.

References

- [1] M. Algrain, J. Quinn, Accelerometer based line-of-sight stabilization approach for pointing and tracking systems, in: Proceedings of Second IEEE Conference on Control Applications, Vancouver, BC, 1993.
- [2] E. Baarsma, M. Collewijn, Vestibulo-ocular and optokinetic reactions to rotations and their interaction in the rabbit, *Journal of Physiology* 238 (1974) 603–625.
- [3] D. Ballard, C.M. Brown, Principles of animate vision, *Computer Vision Graphics and Image Processing* 56 (1) (1992) 3–21.
- [4] B. Barshan, H. Durrant-Whyte, Inertial sensing for mobile robotics, *IEEE Transactions on Robotics and Automation* 11 (1995) 328–342.
- [5] A. Benson, F.E. Guedry, G.M. Jones, Response of semicircular canal dependent units in vestibular nuclei to rotation of a linear acceleration vector without angular acceleration, *Journal of Physiology* 210 (1970) 475–494.
- [6] B. Biguer, C. Prablanc, Modulation of the vestibulo-ocular reflex in eye-head orientation as a function of target distance in man, in: A. Fuchs, W. Becker (Eds.), *Progress in Oculomotor Research*, Elsevier, Amsterdam, 1981.
- [7] R. Brooks, Behavior-based humanoid robotics, in: Proceedings of IEEE/RSJ IROS'96, Vol. 1, 1996, pp. 1–8.
- [8] G.A. Bush, F.A. Miles, Short-latency compensatory eye movements associated with a brief period of free fall, *Experimental Brain Research* 108 (1996) 337–340.
- [9] C. Capurro, F. Panerai, G. Sandini, Dynamic vergence using log-polar images, *International Journal of Computer Vision* 24 (1) (1997) 79–94.
- [10] H. Collewijn, C. Erkelens, Binocular eye movements and the perception of depth, in: E. Kowler (Ed.), *Eye Movements and Their Role in Visual and Cognitive Process: Review of Oculomotor Research*, Elsevier, Amsterdam, pp. 213–261.
- [11] D. Coombs, C. Brown, Intelligent gaze control in binocular vision, in: Proceedings of the Fifth IEEE International Symposium on Intelligent Control, Philadelphia, PA, 1990.
- [12] V. Cornilleau-Pérès, J. Droulez, The visual perception of three-dimensional shape from self-motion and object-motion, *Vision Research* 34 (1994) 2331–2336.
- [13] B.T. Crane, E.S. Viirre, J.L. Demer, The human horizontal vestibulo-ocular reflex during combined linear and angular acceleration, *Experimental Brain Research* 114 (1997) 304–320.

- [14] R.S. Gellman, J.R. Carl, F.A. Miles, Short latency ocular-following responses in man, *Vision of Neuroscience* 5 (1990) 107–122.
- [15] J. Gibson, *The Perception of the Visual World*, Houghton Mifflin, Boston, MA, 1950.
- [16] J.J. Gibson, *The Senses Considered as Perceptual Systems*, Houghton Mifflin, Boston, MA, 1966.
- [17] T. Hine, F. Thorn, Compensatory eye movements during active head rotation for near targets: Effects of imagination, rapid head oscillation and vergence, *Vision Research* 27 (1987) 1639–1657.
- [18] I. Howard, The stability of the visual world, in: F. Miles, J. Wallman (Eds.), *Visual Motion and Its Role in the Stabilization of Gaze*, Elsevier, Amsterdam, pp. 103–118.
- [19] E. Kandel, J. Schwartz, T. Jessel, *Principles of Neuroscience*, Elsevier, Amsterdam, 1991.
- [20] E. Keller, Gain of vestibulo-ocular reflex in the monkey at high rotational frequencies, *Vision Research* 18 (1978) 311–315.
- [21] J. Koenderink, J. van Doorn, Affine structure from motion, *Journal of the Optical Society of America* 8 (2) (1991) 377–385.
- [22] D. Lee, P. Reddish, Plummeting gannets: a paradigm of ecological optics, *Nature* 293 (1981).
- [23] J. Lobo, J. Dias, Towards visual and inertial sensing integration: Contribution to improve mobile robot autonomy, in: *Proceedings of the Fifth Symposium on Intelligent Robotics Systems*, Stockholm, Sweden, 1997.
- [24] J. Micheal, G.M. Jones, Dependence of visual tracking capability upon stimulus predictability, *Vision Research* 16 (1966) 707–716.
- [25] F.A. Miles, K. Kawano, L.M. Optican, Short-latency ocular following responses of monkey. I. Dependence on temporospatial properties of the visual input, *Journal of Neurophysiology* 56 (1986) 1321–1354.
- [26] P. Morasso, G. Sandini, V. Tagliasco, R. Zaccaria, Control strategies in the eye-head coordination system, *IEEE Transactions on Systems, Man and Cybernetics* 7 (9) (1977) 639–651.
- [27] K. Nakayama, Motion parallax sensitivity and space perception, in: A. Hein, M. Jeannerod (Eds.), *Spatially Oriented Behavior*, Springer, New York, pp. 223–241.
- [28] K. Pahlavan, T. Uhlin, J.-O. Eklundh, Integrating primary ocular processes, in: *Proceedings of ECCV92—European Conference of Computer Vision*, Santa Margherita Ligure, Italy, Lecture Notes in Computer Science, Vol. 588, Springer, Berlin, 1992.
- [29] G.D. Paige, Vestibuloocular reflex and its interactions with visual following mechanisms in the squirrel monkey. I. Response characteristics in normal animals, *Journal of Neurophysiology* 49 (1983) 134–168.
- [30] G.D. Paige, Linear vestibulo-ocular reflex (LVOR) and modulation by vergence, *Acta Otolaryngol Supplement* 481 (1991) 282–286.
- [31] F. Panerai, Inertial sensors for controlled camera systems, Technical Report 2/95, LIRA-Lab, DIST University of Genoa—LIRA laboratory, University of Genoa, 1995.
- [32] F. Panerai, Integration of inertial and visual information in binocular vision systems, Ph.D. Thesis, Department of Communication, Computers and Systems Science, University of Genova, Genova, Italy, 1998.
- [33] F. Panerai, G. Sandini, Oculo-motor stabilization reflexes: integration of inertial and visual information, *Neural Networks* 11 (7–8) (1998) 1191–1204.
- [34] P. Questa, G. Sandini, Time to contact computation with a space-variant retina-like c-mos sensor, in: *Proceedings of the International Conference on Intelligent Robots and Systems*, Osaka, Japan, 1996.
- [35] B. Rogers, M. Graham, Similarities between motion parallax and stereopsis in human depth perception, *Vision Research* 22 (1982) 261–270.
- [36] U. Schwarz, C. Busetini, F.A. Miles, Ocular responses to linear motion are inversely proportional to viewing distance, *Science* 245 (1989) 1394–1396.
- [37] U. Schwarz, F.A. Miles, Ocular responses to translation and their dependence on viewing distance. I. Motion of the observer, *Journal of Neurophysiology* 66 (1991) 851–864.
- [38] P. Sharkey, D. Murray, S. Vandeveld, I. Reid, P. McLauchlan, A modular head/eye platform for real-time reactive vision, *Mechatronics* 3 (4) (1993).
- [39] M. Shelhamer, D.M. Merfeld, J.C. Mendoza, Effect of vergence on the gain of the linear vestibulo-ocular reflex, *Acta Otolaryngol (Stockh) Supplement* 520 (1995) 72–76.
- [40] L. Snyder, W. King, Effects of viewing distance and location of the axis of head rotation on the monkey's vestibuloocular reflex. I. Eye movement response, *Journal of Neurophysiology* 67 (1992) 861–874.
- [41] S. Tabak, H. Collewyn, L.J. Boumans, J. Van der Steen, Gain and delay of human vestibulo-ocular reflexes to oscillation and steps of the head by a reactive torque helmet. I. Normal subjects, *Acta Otolaryngol* 117 (1997) 785–795.
- [42] L. Telford, S.H. Seidman, G.D. Paige, Dynamics of squirrel monkey linear vestibuloocular reflex and interactions with fixation distance, *Journal of Neurophysiology* 78 (1997) 1775–1790.
- [43] L. Telford, S.H. Seidman, G.D. Paige, Canal-otolith interactions in the squirrel monkey vestibulo-ocular reflex and the influence of fixation distance, *Experimental Brain Research* 118 (1998) 115–125.
- [44] D. Terzopoulos, T. Rabie, Animat vision: Active vision in artificial animals, *Videre: Journal of Computer Vision Research* 1 (1) (1997) 2–19.
- [45] T. Vieville, O.D. Faugeras, Computation of inertial information on a robot, in: *Proceedings of Fifth International Symposium of Robotic Research*, Tokyo, MIT Press, Cambridge, MA, 1990.
- [46] E. Viirre, D. Tweed, K. Milner, T. Vilis, Reexamination of the gain of the vestibuloocular reflex, *Journal of Neurophysiology* 56 (1986) 439–450.
- [47] C. Weiman, Binocular stereo via log-polar retinas, in: *Proceedings of SPIE AeroSense95*, Orlando, FL, 1995.
- [48] V. Wilson, G.M. Jones, *Mammalian Vestibular Physiology*, Plenum Press, New York, 1979.



Francesco Panerai graduated in Electronic Engineering in 1993 at University of Genova. He did his Ph.D. in Robotics and Computer Vision at LIRA Lab. His doctoral thesis developed the idea of integrating inertial sensors in robot vision systems to improve gaze stabilization in active vision tasks. In 1995 he joined for a semester the Hybrid Vision team at NASA Johnson Space Center, where he worked

on a binocular robot exploiting space-variant sensors and a unique real-time image remapper. In 1998 he started a post-doctoral fellowship at LPPA (Collège de France) working on psychophysics of the visual system. He developed techniques for measuring subjects head movements in 3D space in real-time. These techniques are currently used to investigate the integration of visual and non-visual information in active vision observers.



Giorgio Metta was born in Cagliari, Italy, in 1970. He received the M.S. degree in Electronic Engineering from the University of Genova, Italy, in 1994, discussing a thesis on visual servoing and robot manipulation. In 1995 he was with the Department of Computer Science, Leeds University, England, funded by a grant from the EU-HCM SMART project. After his mandatory national service he did in the

Italian Navy Chief Staff, in 1996 he won a Ph.D. position at the Laboratory for Integrated Advanced Robotics (LIRA Lab, Genova). Currently, he is finalizing his Ph.D. studies. During the three year course he investigated issues related to the development of sensori-motor coordination from both the artificial and the computational neuroscience point of view. For this reasons, he built an anthropomorphic head–arm robotic setup.



Giulio Sandini teaches the course of Natural and Artificial Intelligent Systems for students of the Electronics and Computer Science curriculum. He currently coordinates the activity of researchers at the Laboratory of Integrated Advanced Robotics (LIRA Lab) where research related to robotics and computational neuroscience is carried out. Among the ongoing projects the control of a binocular head using

space-variant, anthropomorphic sensors and the guidance of mobile robots on the basis of visual information. Giulio Sandini has been a member of programme committees of international conferences and chairman and co-chairman of international conferences and workshops.