Two-Mode Control: An Oculomotor-Based Approach to Tracking Systems

Ehud Rivlin, *Member, IEEE*, Héctor Rotstein, and Yehoshua Y. Zeevi

***Abstract*—Control theory principles have been used to gain understanding on how the different components of the oculomotor system operate. The purpose of this paper is to use the knowledge about how the visual system organizes these components to propose a new tracking paradigm.**

The tracking system is assumed to be described by linear time-invariant discrete-time state-space equations. The inputs to this system are the control action and the reference signal. The outputs are the controlled variable z and the measurements. The tracking objective is to keep $|e(t)|$ smaller than some prespecified value γ. This is motivated by the organization of the human visual system, where the target should be kept within the fovea of the eye. The reference signal is modeled as the filtered version of a driving signal $a(t)$. Its characteristics, together with the tracking objective, lead to the design of an optimal controller referred to as the "smooth pursuit controller."

Since the tracking system should continue to operate in the event of a controlled variable constraint violation, a more complex control strategy is required. The approach in this paper, motivated by the behavior of the visual system, is to switch off the smooth controller whenever a violation occurs and design a time-optimal control action, i.e., a "saccade," to drive the control system so that the constraint is satisfied after the shortest possible time interval. After that, the smooth controller is switched back into the loop. The way this switching is performed is critical for obtaining "good behavior"; a method is proposed which is based on a careful definition of the target set for the saccade.

The tracking system proposed in this paper is closely related to recent results in linear optimal and robust control theory. It also shows that the human visual system poses some very interesting questions and open problems which may stimulate further control theoretic work.

***Index Terms*— State-space methods, switching systems, time-optimal control, tracking, visual feedback.**

I. Introduction

CONTROL system theory has had a profound and easily recognizable impact on the investigation and understanding of the oculomotor system [21], [25], [27]. For instance, much of the research in the early 1960's was directed toward finding and analyzing the frequency response of the eye movements when tracking an object. It soon became apparent that a single linear system could not provide a good model, since the tracking of all but very simple movements activates two subsystems, namely the saccadic and the smooth pursuit systems. Therefore, effort was directed toward measuring and characterizing each one of the subsystems/modes independently. The earlier findings and main ideas are summarized by Robinson in a comprehensive tutorial [21].

As simple models failed to predict the behavior of the visual system, more concepts and ideas from control theory were introduced and exploited, and ingenuous devices and experiments were designed to validate or disprove different working hypotheses. Examples are the sample-data model for saccadic control [27], the usage of optimal control theory, and numerical methods to validate the hypothesis that saccadic movements are (approximately) time optimal [9] or the idea of using so-called positive feedback [27] (which turns out to be the familiar Internal Model Principle). An external, secondary feedback loop was also introduced in some experiments to further explore the characteristics of the visual-oculomotor system [29].

By the end of the 1980's, the interest in building vision heads which can alter their position in response to external stimulus gave rise to a renewed interest in the human oculomotor system, with emphasis on using it as a source of control strategies. For example, the notion of using two control laws, one relatively slow and continuous and another fast but discrete in nature, was rapidly adopted into mechanical designs. Most of the resulting schemes were implemented in a largely *ad hoc* manner, relying on lengthy trial-and-error on-line tuning procedures since, except for a few exceptions, control theory failed to address the problem of how to organize the different anthropomorphic-inspired subsystems into an efficient and precise tracking mechanism. This approach worked well for the first prototypes, haunted as they were by mechanical difficulties such as the voluminous size and weight of the cameras. However, current new technologies provide small, lightweight cameras and motors with improved weight/torque characteristics so that reliable vision heads with fast dynamics can be constructed. A systematic approach to the design of the control system has therefore become crucial for achieving the performance that these heads can potentially deliver.

In [24], the authors analyze some of the basic control questions which arise when attempting to implement an anthropomorphic-inspired design for active vision systems. Issues like implementing a foveal window or region of high spatial resolution [19], in a mechanical system based on cameras which have inherent uniform resolution, or using a tracking system based on two modes which must interact to

Manuscript received March 14, 1997. Recommended by Associate Editor, B. K. Ghosh.

E. Rivlin is with the Department of Computer Science, Technion-Israel Institute of Technology, Technion City, Haifa 32000, Israel.

H. Rotstein and Y. Y. Zeevi are with the Department of Electrical Engineering, Technion-Israel Institute of Technology, Technion City, Haifa 32000, Israel.

Publisher Item Identifier S 0018-9286(98)04330-X.

produce improved performance, were addressed in light of novel results in linear optimal control theory. Following this work, it is possible to compute the optimal size of the fovea by trading-off the time delay in the loop and the stringency of the control objective; while the former calls for smaller foveas, the latter is less tight for larger ones. The formulation also produces a linear time-invariant (LTI) controller which can be assimilated to the smooth pursuit mechanism observed in the human eye. Since this controller can only guarantee good performance for a restrictive set of trajectories of the target, a second control which overrides the previous one when specifications fail was also considered; this control law resembles human saccades.

From this research, it became apparent that the solution to the tracking problem provided by the oculomotor system can be put forward as a new approach to tracking. To the best of our knowledge, this has not been attempted before, in spite of the fact that the oculomotor system constitutes a handy and rather successful tracking paradigm. The main purpose of this paper is to show how to organize two relatively simple control laws to achieve improved performance, following a simple model of how the eye tracks. The paper is organized as follows. In Section II, the main characteristics of the oculomotor system and some of the current approaches to tracking are briefly reviewed, and their main differences are highlighted. In Section III the tracking problem is formulated, and smooth pursuit is characterized as a family of linear optimal control problems. This formulation also establishes a sound basis for putting forward a two-mode controller. Section IV is devoted to considering the counterpart of human visual saccades. The main point is not the saccadic control law itself, which is formulated as a problem of time-optimal control to a time-varying set, but rather how to interconnect the two modes without negative transient effects. It is claimed that the approach for tracking is novel, but it has relevance and is related to other areas of control. This is discussed in Section V. Finally, Section VI contains conclusions and an outline of further works.

II. Tracking with the Eye and System Theory

The purpose of this section is to provide some background information. First, the relevant characteristics of the tracking mechanisms in the human visual system are considered. Then, current approaches to tracking systems are briefly reviewed with special emphasis on those dealing with optimal performance. Finally, the main differences between the two approaches are stressed, and conditions under which classical tracking will fail to provide suitable controllers are highlighted.

A. Human Visual System

Most of the processing in human vision is devoted to a very small portion of field of view called "fovea." The foveal field of view is hardly 2 degrees in extent [15], although even within this region there is considerable variation of visual acuity. The movements of the eyes shift the foveal field allowing us high-resolution vision wherever it is needed. Carpenter [8] has classified eye movements into two main categories: gaze-shifting and gaze-holding movements. We are interested especially in two gaze-shifting eye movements, a fast one called *saccades* and a slow one called *smooth pursuit*. The saccades are fast movements under voluntary control which cause a change in the fixation point. The smooth pursuit is a movement which occurs primarily when a moving stimulus is presented on the retina and is part of the mechanisms giving foveal tracking abilities.

Smooth pursuit is effective for relatively slow target trajectories. These movements usually have latency of about 130 ms and velocity of less than 20°/s [15], though they can be accelerated under certain conditions (i.e., using dynamic visual noise [28]) to reach velocities higher than 40°/s. These movements occur when the eye is tracking a smoothly moving target and appears to keep the target image stabilized with respect to the retina. Due to "smaller-than-one" gain, the smooth pursuit does not stabilize the retinal image completely. When there is a considerable deviation of the image of the target from the center of the fovea, the smooth pursuit is interrupted by correcting saccades with a frequency that increases with target velocity [10].

The saccades are executed with high speed (hundreds of °/s) but have latency of as much as 150–250 ms; once triggered, they cannot be influenced by visual information. The time course of saccades to a static object is a nonlinear function of the amplitudes; the larger the movement, the faster the saccade, with saturation at about 1000°/s for a movement of approximately 50° [9]. For some time it was thought that vision is impaired during saccades (the so-called saccadic suppression), but further experiments showed that a certain amount of visual processing occur during the course of a saccade [8].

Tracking of a smoothly moving target in humans is executed by smooth eye movements with position corrections by saccades which compensate for the retinal slip resulting from the "smaller-than-one" gain of the smooth eye movement system. The combination of the two movements results in a rather fast dynamic fixation of the fovea central on the target.

B. A Comparison

From the discussion above, the following differences are to be noted between the oculomotor tracking system and the standard control approach.

- *Structure*: Controllers for linear systems resulting from the tracking approaches discussed above are "simple," in the sense that they are themselves linear; this is justified in many cases (i.e., in linear-quadratic-Gaussian (LQG) or $\mathcal{H}_\infty$) when linear controllers exhibit as a good performance as nonlinear ones. On the other hand, the oculomotor system is composed of two subsystems which interact by means of a switching logic, to achieve a performance objective. Systems of this type are called *hybrid* since they are composed of a dynamical and a discrete-event system. For their analysis, hybrid systems are organized in layers, where in the lower layer the control laws are generated and in the upper a "supervisor" assigns which control law should be generated. Although implicit, this organization will not be mentioned in the sequel.

- *Control Objective*: The main objective of the oculomotor tracking system is to keep the target within the fovea. When this is not possible and the error becomes larger than a fixed threshold, a saccade is triggered that aims to reduce the error to bellow the threshold. The reason for this behavior is that the information obtained from the target decreases drastically when it no longer lies within the fovea, because of the reduced number of sensors. Notice that a high density of sensors produces higher resolution at the cost of larger computational delays, and hence the objective arises as a compromise between these two. As opposed to this, a linear tracking system has a smoother behavior, since its objective is to reduce an overall measurement of performance, like a quadratic norm criterion. It is fair to say that standard tracking approaches sacrifice (or at least de-emphasize) local behavior on behalf of global performance, while the opposite is true for the oculomotor system.

III. Problem Formulation and Smooth Pursuit

Consider an LTI system with a state-space realization

$$\begin{aligned} x(t+1) &= Ax(t) + B_2u(t) + B_1r(t) \\ e(t) &= C_1x(t) + Du(t) - r(t) \\ y(t) &= \begin{bmatrix} x(t) \\ r(t) \end{bmatrix}. \end{aligned} \tag{1}$$

Here $x(t) \in \mathbb{R}^n$ is the state of the system with initial value $x(0) = x_0$, $y(t)$ denotes the available measurements, while $e(t)$ is the difference between the system output $\hat{e}(t) = C_1x(t) + Du(t)$ and the reference signal $r(t)$. For simplicity, both $r(t)$ and $e(t)$ are assumed to be scalar signals. Notice that no process or measurement noises have been included in the model, and the whole state of the plant is assumed to be available for feedback.

The setup for the tracking problem is illustrated in Fig. 1. The controller $\mathcal{K}$ should be designed in such a way that the error $e(t)$ between the reference $r(t)$ and the "output" $\hat{z}(t)$ of the system remains small (in some given sense). As shown in the figure, the controller generates the control action $u(t)$ based on the state vector $x(t)$ of the plant and the reference signal $r(t)$; the controller should be *internally stabilizing*, i.e., so that in the absence of a reference signal any initial state converges to zero. Tracking paradigms differ in how the size of $e(t)$ is measured and in how the reference signal $r(t)$ is characterized, i.e., what class of signals the system should be able to track.

If the human oculomotor system is viewed in this framework, the objective of keeping the target within the fovea can be expressed as having the error $e(t)$ smaller than the half size of the fovea for each time instant *t*; in mathematical terms, this can be expressed as $\|e\|_\infty \leq \gamma$, with $\gamma \approx 1°$. As mentioned in the previous section, experimental work suggests that this is achieved by using one controller designed by penalizing error *velocity* (retinal slip) with another one attempting to reduce large errors (retinal error). These correspond to the smooth pursuit and the saccade subsystems, respectively. More debatable is the characterization of the class of visual signals which can be tracked smoothly (i.e., without resorting to saccadic corrections). For simplicity, it can be assumed that the eyes attempt to track not the position of the target but rather its acceleration. This is suggested, for instance, by the fact that targets with zero acceleration are asymptotically tracked almost independently of their velocity. However, it is not consistent with the mechanism for maintaining fixation, where nonzero steady-state errors may appear; this nonconsistency suggests that the oculomotor system can only partially be explained in terms of simple linear system models (but see the comment below).

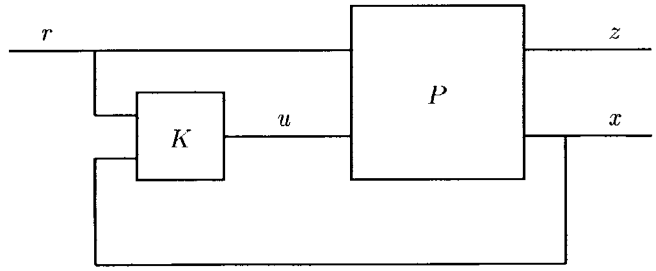

Fig. 1. Configuration for the tracking problem.

For the anthropomorphic-inspired approach to tracking, it will be assumed that the signal $r(t)$ in (1) is generated as the result of passing an unknown signal $a(t)$ through a finite-dimensional LTI "smooth pursuit" filter $\mathcal{F}$

$$r(t) = (\mathcal{F}a)(t). \tag{2}$$

The filter $\mathcal{F}$ has the transfer function $\hat{f}_M(z)$

$$\hat{r}(z) = \hat{f}_M(z)\hat{a}(z).$$

Here the *hatted* notation "$\hat{r}$" denotes the z-transform of $r(t)$, and likewise for other signals. As an example, $\hat{f}_M(z) = 1/[(z-1)^2]$ maps discrete-time accelerations, or more correctly second differences, into positional displacements and induces the asymptotic tracking capability to ramps found in the oculomotor system when tracking moving targets. For reasons to become clear later, the smooth pursuit filter is restricted to be of the form $\hat{f}_M(z) = 1/[d(z)]$, where $d(z)$ is a polynomial of appropriate order. The inclusion of $\mathcal{F}$ implies some *a priori* knowledge of the signals to be tracked. More generally, not a single filter but a collection of them could be used, with a supervising control logic deciding which one should be selected for specific tasks. For instance, one could select $\hat{f}^1_M(z) = 1$ for regulating around zero or following slowly moving targets and $\hat{f}^2_M(z)$ for tracking fast targets. The signal $a(t)$ is assumed to be deterministic and norm-bounded, with the following two different cases of interest:

$$\|a\|_\infty \doteq \sup_t |a(t)|$$

$$\|a\|_2 \doteq \left(\sum_{t=0}^{\infty} |a(t)|^2 \right)^{1/2}.$$

The spaces of sequences which are bounded in the sense of these two norms are called ℓ_∞ and ℓ_2, respectively; the unit balls in these spaces will be as $\mathcal{B}\ell_i$ with $i = 2$ or ∞. Given a controller K which stabilizes the closed-loop system, it is possible to compute the worst case ∞-norm for the output as

$$\begin{aligned}\mu_i(K) &\doteq \sup_{\|a\|_i \le 1} \|z\|_\infty \\ &= \|T_{ea}(K)\|_{\infty, i}.\end{aligned}$$

Here $T_{ea}(K)$ denotes the closed-loop transfer function from a to e and hence $\mu_i(K)$ denotes an induced system norm; internal stability implies that $\mu_i(K) < \infty$.

By linearity, a given controller K can guarantee that the constraint $\|z\|_\infty \le \gamma$ will not be violated if $\|a\|_i \le \epsilon(K) = \gamma/\mu_i(K)$. In order to track the largest possible signals with the norm of the output not exceeding γ, the controller should be selected so that $\mu_i(K)$ can be minimized; any such controller will be called a "smooth pursuit" controller and denoted by K^s. If $\|a(t)\|_\infty$ is assumed to be bounded, then the computation of $\mu_\infty(K^s)$ is an ℓ_1-optimal control problem, for which efficient numerical computed techniques exists (see [12] and also below for additional details). As shown in [13], if there exists a nonlinear time-varying controller meeting the specification for a given ϵ, there also exists an LTI achieving the same performance, and hence one can assume without loss of generality that K^s is LTI. Although similar facts are true for $a(t)$ with bounded two-norm (i.e., the generalized $\mathcal{H}_2$ problem considered in [22]), only the previous problem is considered in the sequel since it simplifies the treatment of saccades.

In summary, there exists an upper bound ϵ over the norm of the signal which can be tolerated while guaranteeing that the constraint over $e(t)$ is not violated. Note that this is a *worst-case* condition: the bound $|e(t)| \le \gamma$ may hold even if $\|a\|_i > \epsilon$. If for some sample instant t_v the constraint on $e(t_v)$ is violated, the controller K^s can no longer function, and a second control law must be generated; this leads naturally to the two-mode tracking system proposed in this paper. It also offers an explanation for its occurrence in the human visual system: if the bound on z is hard, in the sense that its violation produces a drastic reduction of performance, the tracking system should tolerate relatively large errors over a transient period in order to eventually satisfy $|e(t)|_\infty \le \gamma$. Since K^s is optimal, then this can only be achieved by switching to a second controller. This is elaborated in the next section; before that, the parameterization of stabilizing controllers which leads to the computation of K^s is briefly reviewed.

A. Parameterization of Linear Tracking Controllers

Assume that the controller is linear, time-invariant, and finite-dimensional. Then, introducing the parameterization of all stabilizing controllers [30] it is possible to write the transfer function between $\hat{r}$ and $\hat{e}$ as

$$\hat{S}(z) = \hat{S}_1(z) + \hat{S}_2(z)\hat{q}(z) \tag{3}$$

where

$$\begin{aligned}\hat{S}_1(z) &= C_F(zI - A_F)^{-1}B_1 - 1 \\ \hat{S}_2(z) &= C_F(zI - A_F)^{-1}B_2 + D\end{aligned}$$

with

$$A_F = A + B_2F \tag{4}$$
$$C_F = C_1 + DF \tag{5}$$

and F is a matrix such that A_F is stable, i.e., has all its eigenvalues inside the open unit disk. $q(z)$ is the degree of freedom of the parameterization, a stable linear transfer function. Suppose now that $\mathcal{F}$ is equal to the discrete-time double integrator $f_M(z) = 1/[(z-1)^2]$. Then, in order to have a bounded error for bounded acceleration, $\hat{S}(1)$ should be such that $\hat{S}(1) = 0$ and $\lim_{z\to 1} \hat{S}(z)/(z-1) = 0$. By defining $q_0 = -\hat{S}_1(1)/\hat{S}_2(1)$ and $q_1 = -\lim_{z\to 1}((\hat{S}_1(z) + \hat{S}_2(z)q_0)/(z-1))/\hat{S}_2(z)$, the set of all transfer functions resulting from stabilizing controllers, and such that $e(t)$ is bounded for bounded accelerations, may be written as

$$\hat{S}(z) = \hat{S}_a(z) + \hat{S}_b(z)\hat{q}_1(z) \tag{6}$$

where

$$\begin{aligned}S_a(z) &= \hat{S}_1(z) + \hat{S}_2(z)q_0 + S_2(z)\frac{z-1}{s+0.5} \\ \hat{S}_b(z) &= \hat{S}_2(z)\left(\frac{z-1}{z+1}\right)^2\end{aligned}$$

and $\hat{q}_1(z)$ is a stable transfer function. The resulting optimization problems are singular because of the existence of a double zero at one. In the ℓ_1 setting, this type of problem is discussed in [26].

For a generic $\hat{f}_M(z)$, the parameter $\hat{q}(z)$ should be selected in such a way that the *unstable* poles of $\hat{f}_M(z)$ be canceled by zeros of the closed-loop, i.e., for each pole z_i with $|z_i| \ge 1$, one should have $\hat{S}(z_i) = 0$ and corresponding modifications for multiplicities larger than one.

IV. The Saccadic Mode

Motivated by the fact that the resolution in the human eye decreases according to the power low outside the fovea [17], it is assumed in what follows that either the signal-to-noise ratio in $r(t)$ or the calculations involved in processing it increase significantly whenever $|e(t)| > \gamma$. For example, in the active vision system discussed in [24], a foveal window is included to speed up computations. If this window is implemented by spatially down-sampling the signal outside the fovea, then a larger discretization error appears while processing time increases substantially if the fovea is implemented by means of a special purpose camera. As a consequence of this deterioration, the smooth pursuit controller either cannot be used because the information is not available when required, or fails to meet the specifications if the signal becomes noisy. The tracking system must then switch to a second mode, which has the objective of driving the system back to the smooth pursuit regime.

Given an initial time t_0, consider the set bounded signals over a finite interval

$$\mathcal{A}_\infty^\epsilon(t_0, T) \doteq \left\{ a(t) \in \ell_\infty \text{ s.t. } \sup_{t\in[t_0,T]} |a(t)| \le \epsilon \right\} \tag{7}$$

and $\mathcal{A}_\infty^\epsilon(t_0) \doteq \bigcup_{T \ge t_0} \mathcal{A}_\infty^\epsilon(t_0, T)$; in this notation $\mathcal{A}_\infty^1(0) = \mathcal{B}\ell_\infty$. Let $a_r(t) \in \ell_\infty$ be such that $a_r \in \mathcal{A}_\infty^\epsilon(0, t_b) \cap \mathcal{A}_\infty^\epsilon(t_c)$

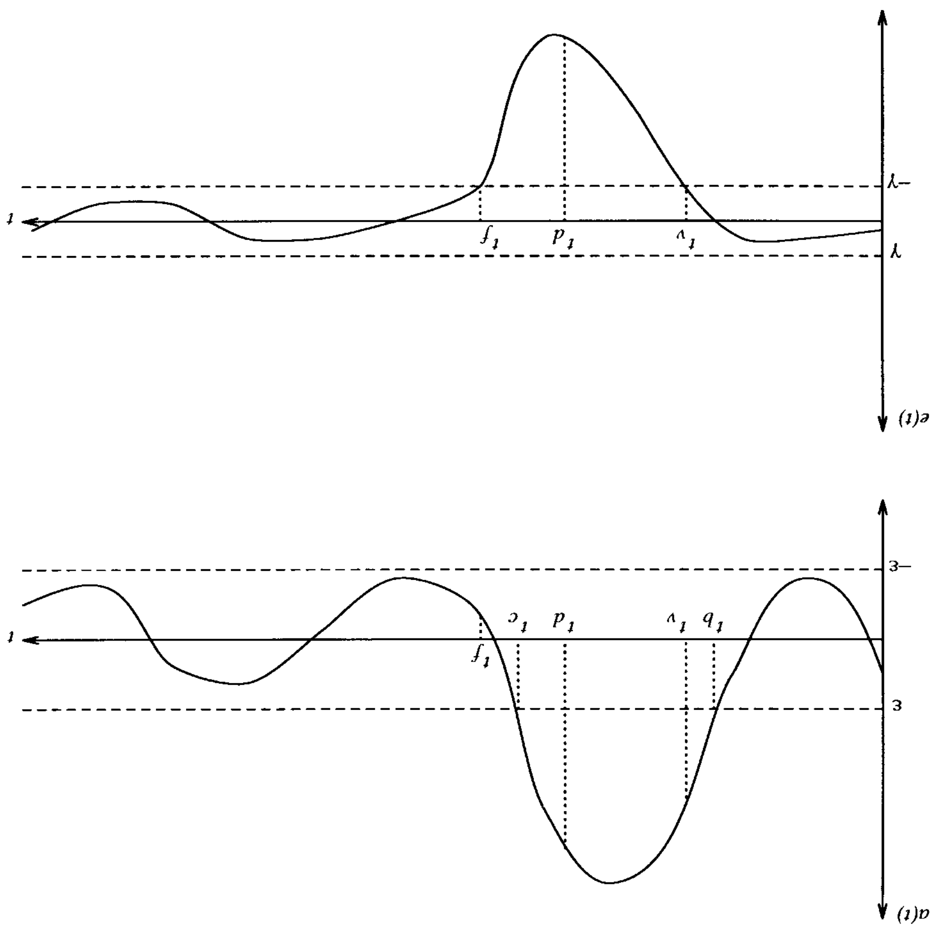

Fig. 2. A signal in $\mathcal{A}_\infty^\epsilon(0, t_b) \cup \mathcal{A}_\infty^\epsilon(t_c)$ and the corresponding error $e(t)$.

for $t_b < t_c$, $a_r \notin \mathcal{A}_\infty^\epsilon(0)$, and for some $t_v \geq t_b$, $|e(t_v)| > \gamma$. The objective of saccadic control is to generate a control action $u^{\mathrm{sac}}(t)$, for $t \in [t_d, t_f]$, where $t_d > t_v$ and $t_f \geq t_c$ so that $|e(t)| \leq \gamma$ can be guaranteed for $t \geq t_f$. Since during the interval $[t_v, t_f]$ the tracking specifications are not achieved, the final time t_f should be made as small as possible. An illustrative example of a signal a_r, together with the corresponding error signal, is shown in Fig. 2.

More generally, $a_r \in \bigcap_j \mathcal{A}_\infty^\epsilon(t_{b_j}, t_{b_{j+1}})$, with $b_j < b_{j+1}$, and the objective is to design $u^{\mathrm{sac}}(t)$, $t \in [t_{d_j}, t_{f_j}]$, with $t_{d_j} > t_{v_j}$ so that $|e(t)| \leq \gamma$ for $t \in [t_{f_j}, t_{v_{j+1}}]$. Since $a_r(t)$, or rather $r(t)$, are not known *a priori*, the sampling instants b_j are unknown and hence u^{sac} can only be computed online. Given a constraint violation at time t_{v_j}, one can attempt to minimize the final time for the saccade t_{f_j}, hoping that $a_r(t) \in \mathcal{A}_\infty^\epsilon(t_{f_j}, T)$ for some T. This is the solution that the oculomotor system seems to have achieved and is the one pursued in the sequel.

The saccadic control $u^{\mathrm{sac}}(t)$ depends on the value of $r(t)$ at some future sample time t_f; since the function is measured online, a model is required to predict this value. The computation of a saccade then involves four distinct stages.

1) *Switch On:* After the constraint on $|e(t)|$ has been violated and before u^{sac} has been computed, the tracking system must be operated "open loop" since the smooth pursuit controller is no longer operational. This is what happens, for instance, to the visual system when a fast moving object is followed and a saccade is triggered by a positional offset. In this stage one can only use the information available at t_v.
2) *Modeling:* The behavior of $r(t)$ must be modeled in order to predict its value at some future sample time. This model should be able to produce accurate predictions of $r(t)$ over a short horizon, but simple enough so that its parameters can be estimated relatively fast.
3) *Saccade:* The control signal u^{sac} must be computed at this stage, based on the value of the state and on the estimate value for the reference signal.
4) *Switch Off:* After the correcting control action has been taken, the smooth controller should be switched back into the loop. It is clear that this switching is critical to guaranteeing the satisfaction of the constraint for $t > t_f$.

Saccadic control can also be formulated in the more general frame of hybrid systems, but the present scheme has the advantage that it can deal effectively with transients by carefully designing the saccadic action. This is discussed next, together with a more detail treatment of the different stages.

A. Switch On

Suppose that the constraint on $|e(t)|$ is violated at time t_v. Then, according to the standing assumption, K^s cannot continue its normal operation and a control action should be computed to meet the specification at some future time. However, this computation takes some time and the tracking system should somehow be driven. A simple solution would be to set $r_v(t) = 0$ for each $t > t_v$, but this can lead to poor performance. Instead, one could select r_v in such a way that the error criterion remains constant; the resulting virtual reference signal may be computed as

$$r_v(t) = C(x(t) - x(t_v)) + D(u(t) - u(t_v)) + r(t_v). \tag{8}$$

Several other choices are possible, depending on the particular application and the *a priori* knowledge available on r.

B. Modeling

Before the saccadic control action is designed, it is necessary to identify a model that can accurately predict $r(t)$ for $t \geq t_d$. This model is referred to in the sequel as the "saccadic model" and is not to be confused with the reference model introduced in the previous section.

This stage is the one called "tracking" (e.g., in [2]), which contains an array of different algorithms to achieve the goal. The specific algorithm should be selected depending on the available *a priori* knowledge on the reference signal; this selection is important since it will determine the number of sample instants required for having an accurate prediction of $r(t)$. In the active vision field, $\alpha - \beta$ or $\alpha - \beta - \gamma$ filters are used because of their simplicity, with their coefficients selected by using the steady-state solution of a Kalman filtering problem [2].

C. Saccade

This is the main step. It is assumed in what follows that the smooth controller K^s has been implemented as

$$u = K^s \begin{bmatrix} x \\ r \end{bmatrix} = Fx + Q^s r$$

where F is a matrix such that A_F as in (4) is stable and Q^s is a stable transfer function. From the parameterization of all stabilizing controllers [30], it is always possible to implement K^s in such a way. Moreover, it is assumed that Q^s (and hence K^s) is *real-rational*, i.e., it has a state-space representation. Consider then the following minimal state-space realization for Q^s:

$$\begin{aligned} x_Q(t+1) &= A_Q x_Q(t) + B_Q r(t) \\ u_1(t) &= C_Q x_Q(t) + D_Q r(t) \end{aligned}$$

with A_Q stable (i.e., all the eigenvalues of A_Q are in the open unit disk). The system from r to e may be represented as

$$\begin{aligned} \begin{bmatrix} x(t+1) \\ x_Q(t+1) \end{bmatrix} &= \begin{bmatrix} A_F & B_2 C_Q \\ 0 & A_Q \end{bmatrix} \begin{bmatrix} x(t) \\ x_Q(t) \end{bmatrix} \\ &\quad + \begin{bmatrix} B_2 + B_1 D_Q \\ B_Q \end{bmatrix} r(t) \\ e(t) &= C_F x(t) + D C_Q x_Q(t) + (D D_Q - 1) r(t). \end{aligned}$$

Consider now the smooth pursuit filter

$$\hat{f}_M(z) = \frac{1}{z^\nu + \alpha_1 z^{\nu-1} + \alpha_2 z^{\nu-2} + \cdots + \alpha_\nu}$$

with state-space realization

$$\begin{aligned} x_M(t+1) &= A_M x_M(t) + B_M a(t) \\ r(t) &= C_M x_M(t) \end{aligned} \tag{9}$$

where

$$\begin{aligned} A_M &= \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ -\alpha_\nu & -\alpha_{\nu-1} & -\alpha_{\nu-2} & \cdots & -\alpha_1 \end{bmatrix} \\ B_M &= \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix} \\ C_M &= [1 \quad 0 \quad \cdots \quad 0]. \end{aligned}$$

The filter $\hat{f}_M(z)$ is assumed to be *anti-stable* so that A_M has all its eigenvalues outside the unit disk (i.e., if λ_i is an eigenvalue of A_M, then $|\lambda_i| \geq 1$); this is without loss of generality since all stable dynamics may be absorbed into the description of the plant above. Let $x_{CL} = [x^T \quad x_Q^T \quad x_M^T]^T$ denote the closed-loop state vector. The closed-loop equations have the form

$$\begin{aligned} x_{CL}(t+1) &= A_{CL} x_{CL}(t) + B_{CL} a(t) \\ z(t) &= C_{CL} x_{CL}(t) \end{aligned}$$

where

$$\begin{aligned} A_{CL} &= \begin{bmatrix} A_F & B_2 C_Q & (B_1 + B_2 D_Q) C_M \\ 0 & A_Q & B_Q C_M \\ 0 & 0 & A_M \end{bmatrix} \\ B_{CL} &= \begin{bmatrix} 0 \\ 0 \\ B_M \end{bmatrix} \\ C_{CL} &= [C_F \quad D C_Q \quad (D D_Q - 1) C_M]. \end{aligned}$$

Assuming $x_{CL}(0) = 0$, let $\mathcal{T}$ denote the LTI operator that maps $a \in \mathcal{A}_\infty^\epsilon(0, T)$ into closed-loop state vector

$$x_{CL}(T) = (\mathcal{T} a)(T). \tag{10}$$

Since K^s stabilizes the closed loop, the system from a to e is stable and hence the modes x_M are all unobservable. Taking $x_a = x + X_1 x_M$, $x_b = x_Q + X_2 x_M$, where the matrices X_1, X_2 solve the Sylvester equation

$$\begin{bmatrix} (B_1 + B_2 D_Q) C_M \\ B_Q C_M \end{bmatrix} + \begin{bmatrix} X_1 \\ X_2 \end{bmatrix} A_M - \begin{bmatrix} A_F & B_2 C_Q \\ 0 & A_Q \end{bmatrix} \begin{bmatrix} X_1 \\ X_2 \end{bmatrix} = 0$$

exposes the unobservable states, and the system from a to z may be written as

$$\begin{aligned} \begin{bmatrix} x_a(t+1) \\ x_b(t+1) \\ x_M(t+1) \end{bmatrix} &= \begin{bmatrix} A_F & B_2 C_Q & 0 \\ 0 & A_Q & 0 \\ 0 & 0 & A_M \end{bmatrix} \begin{bmatrix} x_a(t) \\ x_b(t) \\ x_M(t) \end{bmatrix} \\ &\quad + \begin{bmatrix} X_1 B_M \\ X_2 B_M \\ B_M \end{bmatrix} a(t) \end{aligned} \tag{11}$$

$$e(t) = [C_F \quad D C_Q \quad 0] \begin{bmatrix} x_a(t) \\ x_b(t) \\ x_M(t) \end{bmatrix} \tag{12}$$

or, after eliminating the unobservable states

$$\begin{bmatrix}\dot{x}_a(t)\\ \dot{x}_b(t)\end{bmatrix} = \begin{bmatrix}A_F & B_2C_Q\\ 0 & A_Q\end{bmatrix}\begin{bmatrix}x_a(t)\\ x_b(t)\end{bmatrix} + \begin{bmatrix}X_1B_M\\ X_2B_M\end{bmatrix}a(t) \tag{13}$$

$$e(t) = [C_F \quad DC_Q]\begin{bmatrix}x_a(t)\\ x_b(t)\end{bmatrix} \tag{14}$$

with $x_a(0) = x(0) + X_1x_M(0)$, $x_b(0) = x_Q(0) + X_2x_M(0)$; $x_r = [x_a^T \;\; x_b^T]^T$ is referred to as the *reduced state vector*. The new representation is internally stable and will be assumed to be minimal for simplicity. Let $\mathcal{T}_r$ denote the linear time invariant operator that maps signals $a \in \mathcal{A}_\infty^\epsilon(0, T)$ into the reduced state-vector at time T

$$\begin{bmatrix}x_a(T)\\ x_b(T)\end{bmatrix} = (\mathcal{T}_r a)(T)$$

and let $\mathcal{X}$ be the set of *reachable* states

$$\mathcal{X} \doteq \left\{\begin{bmatrix}x_1\\ x_2\end{bmatrix} = (\mathcal{T}_r a)(T), \text{ for some } T,\, a \in \mathcal{A}_T\right\} \tag{15}$$

where $\mathcal{X}$ is closed, bounded, and is said to be *Positive Invariant* [3], since $[x_a(T)^T \quad x_b(T)^T]^T \in \mathcal{X}$ at T implies $\left[\begin{smallmatrix}x_a(t)\\ x_b(t)\end{smallmatrix}\right] \in \mathcal{X}$ for all $t > T$, if $a \in \mathcal{A}_\infty^\epsilon(T)$. Since $\mathcal{X}$ is associated with a smooth controller K^s, $x \in \mathcal{X}$ implies that $|e(t)| \leq \gamma$ for the corresponding $e(t)$. This fact has far-reaching consequences, as illustrated by the work of Blanchini and Sznaier [4] on static-state feedback for ℓ_1-optimal control.

The key observation that will allow the construction of a target set for saccades is that if the plant is driven so that for some time t_f, the reduced state $x_r \in \mathcal{X}$, then the corresponding control action satisfies the specifications of saccadic control. This is because the state vector of the controller, stored in the computer used for control, can be initialized arbitrarily; hence only x_a is constrained. As it turns out, this can also be used for solving a more general static feedback ℓ_1 problem [23].

Using the saccadic model for the signal $r(t)$ constructed in the modeling stage, compute $\overline{r}(t_f), \overline{r}(t_f+1), \cdots, \overline{r}(t_f + \nu - 1)$, where t_f is some future time and the notation $\overline{r}$ is used to denote that not the true values but some estimates are available. By the state-space model in (3)

$$\overline{x}_M(t_f) = \begin{bmatrix}\overline{r}(t_f)\\ \overline{r}(t_f+1)\\ \vdots\\ \overline{r}(t_f+\nu-1)\end{bmatrix}.$$

The state vector for the reference model at time t_f can then be reconstructed based only on the measurement from t_f to $t_f+\nu-1$, i.e., independently of the input over the same period. This is possible due to the special structure assumed for the smooth pursuit filter, which plays the role of the full state measurability assumed for the plant. Indeed, since the filter is an artifice for formulating the smooth control problem, its (fictitious) state vector cannot be measured, but the structure of the filter is such that the state vector can be determined in a unique way from the signal $r(t)$, which can be measured or, in the current case, estimated.

Given a vector x_M^0, consider the set $\mathcal{E}(x_M^0, \tau)$ of all signals that drive the state vector of the saccadic filter from zero to x_M^0 at some time $\tau > 0$

$$\mathcal{E}(x_M^0, \tau) \doteq \left\{a_v \in \mathcal{A}_\infty^\epsilon(0,\tau) \text{ such that } x_M^0 = \sum_{i=1}^{\tau} A_M^{i-1}B_M a_v(\tau - i)\right\}. \tag{16}$$

Equivalently, $a_v \in \mathcal{E}(x_M^0, \tau)$ if we have (17), as shown at the bottom of the page. It is easy to see that $\mathcal{E}(x_M^0, \tau) \subset \mathcal{E}(x_M^0, \tau+1)$; take $\mathcal{E}(x_M^0) = \bigcup_{\tau\geq 0}\mathcal{E}(x_M^0, \tau)$. Since A_M is assumed to have its eigenvalues outside the open unit disk, τ cannot be taken large due to the numerical problems associated with taking powers of the matrix and verifying the equality (17). From a computational point of view, one should then consider an upper bound $\overline{\tau}$ for τ.

Given a state vector for the smooth pursuit filter x_M^0, consider the set

$$\begin{aligned}\mathcal{O}(x_M^0) &\doteq [I \quad 0 \quad 0]\mathcal{T}\mathcal{E}(x_M^0)\\ &= \{x \text{ such that } x = [I \quad 0 \quad 0](\mathcal{T}a_v)(\tau),\\ &\qquad a_v \in \mathcal{E}(x_M^0, \tau) \text{ for some } \tau > 0\}\end{aligned} \tag{18}$$

$$\begin{aligned}&x_M^0 = [A_M^{\tau-1}B_M \quad A_M^{\tau-2}B_M \;\cdots\; A_MB_M \quad B_M]\begin{bmatrix}a_v(0)\\ a_v(1)\\ a_v(2)\\ \vdots\\ a_v(\tau-1)\end{bmatrix}\\ &|a_v(i)| \leq \epsilon, \qquad 0 \leq i \leq \tau - 1\end{aligned} \tag{17}$$

$$\begin{bmatrix}x(t_f)\\ x_Q^o\\ x_M(t_f)\end{bmatrix} = [A_{CL}^{\tau-1}B_{CL} \quad A_{CL}^{\tau-2}B_{CL} \;\cdots\; A_{CL}B_{CL} \quad B_{CL}]\begin{bmatrix}a_v(0)\\ a_v(1)\\ a_v(2)\\ \vdots\\ a_v(\tau-1)\end{bmatrix}$$

where $\mathcal{T}$ is as defined in (10). Equivalently, $x \in \mathcal{O}(x_M^0)$ if for some $\tau > 0$

$$\begin{bmatrix} x \\ x_M^0 \end{bmatrix} = \begin{bmatrix} I & 0 & 0 \\ 0 & 0 & I \end{bmatrix} \cdot [A_{CL}^{\tau-1}B_{CL} \quad A_{CL}^{\tau-2}B_{CL} \cdots A_{CL}B_{CL} \quad B_{CL}] \cdot \begin{bmatrix} a_v(0) \\ a_v(1) \\ a_v(2) \\ \vdots \\ a_v(\tau-1) \end{bmatrix}$$
$$|a_v(i)| \leq \epsilon,\ 0 \leq i \leq \tau - 1.$$

By the same numerical considerations as before, it is preferable to work with the smaller set $\mathcal{O}(x_M^0)_{\overline{\tau}} \doteq [I \quad 0 \quad 0]\mathcal{T}\mathcal{E}(x_M^0, \overline{\tau})$, which can be described by a finite number of equalities and inequalities. The main result of this section can now be stated.

Theorem 1: Assume that $a_r \in \mathcal{A}_\infty^\epsilon(0, t_b) \cap \mathcal{A}_\infty^\epsilon(t_c)$ for $t_b < t_c$, $a_r \notin \mathcal{A}_\infty^\epsilon(0)$, and for some $t_v > t_b$, $|e(t_v)| > \gamma$ (refer again to Fig. 2). Let $x(t_d)$ be the state vector of the plant at time t_d, $x_M(t_f) \in \mathcal{O}[x_M(t_f)]$ be the state vector of the smooth pursuit filter at time t_f, and let $u^{\mathrm{sac}}(t)$, $t_d \leq t \leq t_f$ be such that

$$x(t_f) = A^{t_f-t_d}x(t_d) + \sum_{i=1}^{t_f-t_d} A^{i-1}B_2u^{\mathrm{sac}}(t_d+i-1) \in \mathcal{O}[x_M(t_f)]. \tag{19}$$

It is then possible to find a state vector $x_Q(t_f)$ for Q^s, so that if the smooth pursuit controller is used for $t \geq t_f$, $|e(t)| \leq \gamma$

Proof: By the assumption that $x(t_f) \in \mathcal{O}(x_M(t_f))$, there exists $\tau > 0$ and a signal $a_v \in \mathcal{A}_\infty^\epsilon(0, \tau)$ such as that shown in the at the bottom of the previous page, for some vector x_Q^o. Define the signal

$$a_{vr} \doteq \begin{cases} a_v(t) & 0 \leq t < t_f \\ a_r(t) & t_f \leq t \end{cases}.$$

From the choice of K^s, it follows that $|(\mathcal{T}a_{vr})(t)| \leq \gamma$ for each $t \geq 0$.

Assume now that the plant is driven by using u^{sac} to $x(t_f)$ at time t_f, and set $x_Q(t_f) = x_Q^o$. Since the resulting closed-loop state vector coincides with $(\mathcal{T}a_{vr})(t_f)$ and $a_r(t) = a_{vr}(t)$, $t \geq t_f$, it follows from basic state-space theory that $(\mathcal{T}a_r)(t) = (\mathcal{T}a_{vr})(t)$ for $t \geq t_f$ and the proof follows. □

The state $x(t_d)$ results from driving the plant from $t = t_v$ to $t = t_d$ with the input signal $r^v(t)$, and hence can be computed for any $t \in [t_v, t_d]$. The lapse between t_v and t_d is used to identify a saccadic filter and then compute the saccadic control law based on the predictions of this filter for the signal $r(t)$. This computation may be formulated as the optimization problem as shown in (20) at the bottom of the page. In order to get the optimal solution, the limit $\tau \to \infty$ should be taken; by the previous numerical considerations, it is convenient to replace τ by a finite upper bound $\overline{\tau}$. Notice that for each fixed t_f, the problem reduces to finding a feasible solution for a linear problem, which can be done efficiently so that an optimal (or $\overline{\tau}$-suboptimal) solution may be computed iteratively as in [14]. Additional constraints (e.g., on the control authority) may also be added in the formulation. The sample instant t_d is required *a priori* for computing u^{sac} and is bounded below by the time it takes to identify the saccadic filter, plus an upper bound on the computation time for u^{sac}.

The fact that the driving signal is constrained to lie in ℓ_∞ simplifies both the derivations and interpretations of the results discussed above. However, it is also possible to consider other norms, in particular the ℓ_2 one, along the lines discussed in [23].

V. Connections with Other Works

The two-mode tracking paradigm presented in this paper has been motivated by the way the human visual system organizes the different components of the oculomotor system. The approach has connections with other works on active vision tracking and linear control theory, which are briefly revised next.

The need for better control for active vision systems constitutes the original motivation for the current work. Until recently, tracking in computer vision was done mostly with a static camera and hence was more related to tracking as understood in [2], both in two-dimensional (2-D), i.e., tracking of an image, and in three-dimensional (3-D). In this context, it is very common to find the usage of Kalman filtering as an off-the-shelf estimator. A survey of these methods and algorithms for the case of a fixed camera can be found in

$$\min \quad t_f \qquad \text{such that:}$$
$$x(t_f) = A^{t_f-t_d}x(t_d) + \sum_{i=1}^{t_f-t_d} A^{i-1}B_2u^{\mathrm{sac}}(t_d+i-1)$$
$$\begin{bmatrix} x(t_f) \\ \overline{r}(t_f) \\ \overline{r}(t_f+1) \\ \vdots \\ \overline{r}(t_f+\nu-1) \end{bmatrix} = \begin{bmatrix} I & 0 & 0 \\ 0 & 0 & I \end{bmatrix} [A_{CL}^{\tau-1}B_{CL} \cdots A_{CL}B_{CL} \quad B_{CL}] \begin{bmatrix} a_v(0) \\ a_v(1) \\ a_v(2) \\ \vdots \\ \overline{r}(t_f+\nu-1) \\ a_v(\tau-1) \end{bmatrix}$$
$$|a_v(i)| \leq \epsilon, \qquad 0 \leq i \leq \tau - 1 \tag{20}$$

[1]. When cameras become mobile, and motors are included in the loop, then the problem becomes one of tracking in the sense described in this paper. In [6] and [7] tracking of the Rochester Robot head is described. Tracking was done by using a proportional–integral–derivative (PID) controller for the camera, driven by the retinal positional error of the image of the target in the dominant camera. In [5] and [11] tracking of the Oxford head was described. The tracking makes use of image motion and position rather than position alone. Using this head to track corner clusters was described in [20]. Corners are tracked from frame to frame using a constant image velocity Kalman filter.

The research in active vision tracking borrows heavily from the work of Bar-Shalom and coworkers [2], where tracking is understood in a sense closely related to estimation or filtering. Some connections may be recognized between the present work and the problem of maneuvering target presented there, in which case the target is modeled by the discrete-time dynamics

$$x_{k+1} = F_k x_k + G_k u_k + v_k \tag{21}$$

where v_k is a zero-mean, white random noise, and u_k is the unknown input to the system; a maneuver then corresponds to a relatively large excursion of the input signal. Two broad approaches to the problem are offered according to whether u_k is assumed to be: 1) a random process or 2) nonrandom and estimated in real-time. These approaches are *multimode*: either several noise levels are tried, several models are run in parallel, and compared, or filters with different degrees are used for the "normal" and "maneuvering" operations. A statistic criterion, like maximum likelihood, is used to decide which output should be selected.

The tracking notions discussed are also related to recent results in linear control. In particular, the optimal smooth controller is constructed as an L_1 or a generalized $\mathcal{H}_2$-optimal control [4], [22], while the mechanism for switching is inspired by [4]. The fact that a switching control law is designed also suggests connections with adaptive control, in particular, with adaptive stabilization with relaxed assumptions and switching-based control laws [16], [18]. The setting there, though, is quite different: a number of controllers are assumed to be known in advance and are switched on and off the loop according to some high-level logic. The motivation for this problem comes from robust control: it is assumed that the "true" plant one wants to control lies on a relatively large set, which cannot be properly controlled by a single LTI controller. The critical problem of initializing the controller after each switching was solved in [18] by assuming that all the controllers share a state space and differ only on the output matrix. Such a solution is not possible in the approach introduced in this paper since the controllers are not known in advanced but rather a control action is computed based on the reference signal and is nonlinear.

VI. Conclusions and Further Work

In this paper, a two-mode approach to tracking motivated by the human oculomotor system has been presented. Following the simple model also considered in [24], visual tracking is based upon two distinct control laws, which appear to have evolved as a tradeoff between computational delays and tightness of the control specifications. More generally, a two-mode mechanism seems to be beneficial whenever the resources available for control are in some sense constrained (for instance, if there are limitations on the computational resources or times) or hard specifications need to be met "most" of the time.

The tracking objective considered is to keep the value of the controlled variable below some prespecified bound. In order to achieve this objective, a smooth controller is designed, assuming that the reference signal $r(t)$ is generated by a signal $a(t)$ passing through a filter. If the driving signal belongs to some normed space and, moreover, is within a ball of radius ϵ, then a controller can be chosen optimally by maximizing the tolerance ϵ. In the event that the error signal violates the constraint, the smooth controller has to be replaced by a control strategy which attempts to drive the system back to specifications in the shortest possible time. This control action is chosen so as to provide an adequate switching between the two modes of operation. Following the anthropomorphic paradigm, the first mode is called "smooth pursuit" and the second "saccadic control."

It is worth stressing that more work should be done to implement the two-mode tracking in practical applications. If only the smooth controller is taken into account, then this transition can be done with relative ease, since the problem reduces to a standard one in robust optimal control. It is then possible to modify the formulation so that a more realistic situation in which: 1) some outputs and not all the state vector are available for feedback; 2) signals are corrupted by noise; and 3) norm bounded plant uncertainty is considered. If noise and uncertainty are appropriately characterized, then a controller can be computed by using established techniques.

On the other hand, several issues should be clarified in order to implement the saccadic control strategy, since even in the idealized situation considered in this paper, the computational cost appears to be very high. The current approach depends on the fact that all states are available for feedback, and it is not clear at this point how this condition may be relaxed if the smooth controller is designed following an L_1 criterion. Moreover, the behavior of the time-optimal control and subsequent switching to the smooth controller deserves more study if noise and plant uncertainty are to be incorporated into the picture. These observations pose some interesting and challenging control problems, some of which are currently under study (see, e.g., [4]).

Currently, the research advances along the following three avenues. First, a tracking system designed according to the approach described in this paper is being implemented on the Technion Robot Head. Second, some of the claims regarding the oculomotor system, like the optimality of smooth pursuit, are planned to be contrasted with actual experimental data. Third, the theoretical control problems are pursued; preliminary results suggest that the computational effort may be lowered if the optimality criterion is replaced by a suboptimal one. It is expected that results that will emerge from this research will both advance the understanding of the human visual tracking and will give rise to better active vision and general tracking systems.

REFERENCES

[1] M. Andersson, "Tracking methods in computer vision," Ph.D. dissertation, Kungl Tenkinska Hogskolan (KTH), Stockholm, Sweden, 1994.
[2] Y. Bar-Shalom and T. E. Fortmann, "Tracking and data association," in *Mathematics in Science and Engineering.* New York: Academic, 1988.
[3] F. Blanchini, "Constrained control for perturbed linear systems," *Control and Dynamical Systems,* vol. 51, C. T. Leondes, Ed. New York: Academic, 1992, pp. 129–182.
[4] F. Blanchini and M. Sznaier, "Persistent disturbance rejection via static state feedback," to be published.
[5] K. Bradshow, P. McLauchlan, I. Reid, and D. Murray, "Saccade and pursuit on an active head/eye platform," *Image and Vision Computing,* pp. 155–159, 1994.
[6] C. Brown, "Gaze controls with interactions and delays," *IEEE Trans. Syst., Man, Cybern.,* vol. 20, pp. 518–527, 1990.
[7] ——, "Prediction and cooperation in gaze control," *Biolobical Cybern.,* vol. 63, pp. 61–70, 1990.
[8] R. H. S. Carpenter, *Movements of the Eyes.* Pion, 1988.
[9] M. R. Clark and L. Stark, "Time optimal behavior of human saccadic eye movement," *IEEE Trans. Automat. Contr.,* vol. 20, pp. 345–348, 1975.
[10] H. Collewijn and E. Tamminga, "Human smooth and saccadic eye movements during voluntary pursuit of different target motions on different backgrounds," *J. Physiology,* vol. 351, pp. 217–250, 1984.
[11] I. D. R. D. W. Murray, P. F. McLauchlan, and P. M. Sharkey, "Reactions to peripheral image motion using a head/eye platform," in *Proc. Int. Conf. Computer Vision,* 1993, pp. 403–411.
[12] M. Dahleh and J. B. Pearson, "l^1-optimal feedback controllers for MIMO discrete-time systems," *IEEE Trans. Automat. Contr.,* vol. 32, no. 4, pp. 314–322, 1987.
[13] M. A. Dahleh and J. S. Shamma, "Rejection of persistent bounded disturbances: Nonlinear controllers," *Syst. Contr. Lett.,* vol. 18, pp. 245–252, 1992.
[14] J. H. de Vlieger, H. B. Verbruggen, and P. M. Bruijn, "A time-optimal control algorithm for digital computer control," *Automatica,* vol. 18, no. 2, pp. 239–244, 1982.
[15] R. W. Dithcburn, *Eye-Movements and Visual Perception.* Oxford, U.K.: Clarendon, 1973.
[16] M. Fu and B. R. Barmish, "Adaptive stabilization of linear systems via switching control," *IEEE Trans. Automat. Contr.,* vol. 31, pp. 1097–1103, 1986.
[17] R. E. Kronauer and Y. Y. Zeevi, "Reorganization and diversification of signals in vision," *IEEE Trans. Syst., Man, Cybern.,* vol. 15, pp. 91–101, 1985.
[18] A. S. Morse, "Supervisory control of families of linear set-point controllers—Part 1: Exact matching," to be published.
[19] N. Peterfreund and Y. Y. Zeevi, "Nonuniform image representation in area-of-interest systems," *IEEE Trans. Image Processing,* to be published.
[20] I. D. Reid and D. W. Murray, "Tracking foveated corner clusters using affine structure," in *Proc. Int. Conf. Computer Vision,* 1993, pp. 76–83.
[21] D. A. Robinson, "The oculomotor control system: A review," *Proc. IEEE,* vol. 56, pp. 1032–1049, 1968.
[22] M. Rotea, "The generalized $\mathcal{H}_2$ control problem," *Automatica,* vol. 29, no. 2, pp. 373–385, 1993.
[23] H. Rotstein and L. Mirkin, "On static feedback for the $\mathcal{L}_1$ and other optimal control problems," to be published.
[24] H. Rotstein and E. Rivlin, "Control of a camera for active vision: Foveal vision, smooth tracking and saccade," to be published.
[25] L. Stark, *Neurological Control Systems.* New York: Plenum, 1968.
[26] M. Vidyasagar, "Further results on the optimal rejection of persistent bounded disturbances," *IEEE Trans. Automat. Contr.,* vol. 36, pp. 642–652, 1991.
[27] L. R. Young, J. D. Forster, and N. van Houtte, "Pursuit eye tracking movements," in *Control of Eye Movements,* C. C. P. Bach-y Rita and J. Hyde, Eds. New York: Academic, 1971.
[28] Y. Y. Zeevi and A. Medina, "Acceleration perceived with dynamic visual noise," *J. Optical Soc. Amer.,* vol. 1, pp. 562–564, 1984.
[29] Y. Y. Zeevi, E. Peli, and L. Stark, "Study of eccentric fixation with secondary visual feedback," *J. Optical Soc. Amer.,* vol. 69, pp. 669–675, 1979.
[30] K. Zhou, J. Doyle, and K. Glover, *Robust and Optimal Control.* Englewood Cliffs, NJ: Prentice-Hall, 1995.

Ehud Rivlin (S'90–M'95) received the B.Sc. and M.Sc. degrees in computer science, the M.B.A. degree from the Hebrew University, Jerusalem, and the Ph.D. degree from the University of Maryland, College Park.

Currently he is an Assistant Professor in the Computer Science Department at the Technion-Israel Institute of Technology. His research interests include machine vision and robot navigation.

Héctor Rotstein received the Ingeniero Electricista degree from the Universidad Nacional del Sur, Argentina, in 1985 and the M.Sc. and Ph.D. degrees in electrical engineering from the California Institute of Technology, Pasadena, in 1990 and 1993, respectively.

From 1985 to 1989, he was employed by the Planta Piloto de Ingenieria Quimica, Bahia Blanca, Argentina, as an Engineer in the Control Group. From 1993 to 1997, he was with the Department of Electrical Engineering of the Technion-Israel Institute of Technology. Since 1997, he has been with Rafael as a Research Engineer. His current research interests include robust optimal control, the control of active vision systems, and the application of system theory to model autonomous systems.

Yehoshua Y. Zeevi (M'90) is a graduate of Technion-Israel Institute of Technology, received the M.Sc. degree from the University of Rochester, NY, and received the Ph.D. degree from the University of California, Berkeley.

He is the Barbara and Norman Seiden Professor of Computer Sciences and the Dean of the Faculty of Electrical Engineering, Technion-Israel Institute of Technology. He was a Vinton Hayes Fellow at Harvard University and has been a regular visitor there. He was also a Visiting Professor at MIT and the CAIP Center of Rutgers University. His research interests include image representation, visual communication, and biological and computer vision. He is one of the founders of iSight, Inc, a company devoted to digital video cameras that mimic visual systems and to real-time image processing, and of UltraGuide, Ltd., a medical technology company that develops computerized systems for visual guidance in ultrasound-assisted procedures.

Dr. Zeevi is the Editor-in-Chief of the *Journal of Visual Communication and Image Representation* and the member of the editorial boards of several other journals.