Optimal Control Theory of Normal and Pathological Slow Eye Movements

Christopher M. Harris[*1], Jonathan Waddington[2]

[1,2] Centre for Robotics and Neural Systems

University of Plymouth, Plymouth, UK

cmharris@plymouth.ac.uk

***Abstract*- Biological eye movements serve vision by controlling the orientation of the retinas. They are under adaptive control implying a control objective, which we frame as minimising a "visual Lagrangian" that does not depend on motor control. Using optimal control theory, we show that the global optimum can be reached if the motor plant is linear and the zero is cancelled (singular control). It appears that the zero needs to be cancelled before the full Lagrangian can be optimised, implying the need for sequential adaptive controllers. We apply this theory to infantile nystagmus syndrome to argue that oscillatory eye movements are optimal, but the waveform depends on whether or not the zero is cancelled. We also show that the role of saccades is crucial in determining the boundary conditions and hence the local optima (fields of extremals). The local optima could be reached but saccades need to change eye velocity as well as eye position. We conclude that optimal control theory can be applied usefully to understanding adaptive biological processes without the need for detailed knowledge of the adaptive control circuits.**

Keywords- Optimization; Optimal Control; Oculomotor; Calculus of Variations; Adaptive Control; Nystagmus; Fixation; Smooth Pursuit; Saccade

I. INTRODUCTION

Since the 1960's, control theory has played a central role in our understanding of biological eye movements. The seminal characterization of the oculomotor plant as a 4th order linear system by David Robinson [1] led to numerous physiological experiments that have culminated in a detailed description of the primate eye movement system.

It is now widely accepted that oculomotor control evolved around a quasi-linear servo system that minimises retinal image motion (retinal slip) caused by perturbation of the head during motion, via the vestibulo-ocular reflex (VOR) and optokinetic sub-systems (OKR). Robinson's plant model is typically approximated by a 2nd order system (an additional complex pole pair [1] is assumed negligible for most purposes). The two poles and a zero are cancelled by a pre-motor compensator, so that the eye closely follows the velocity command (Fig.1) [2]. With the evolution of foveal vision came the need for precise positional control, and the original system was supplemented with the smooth pursuit, fixation, and vergence subsystems to allow foveal tracking of stationary and moving objects in 3 dimensional space. Positional errors are corrected by the saccadic system which generates velocity pulses that rapidly shift eye position [3], but the long delays in the visual system prevent the direct use of negative feedback. Instead, the system generates internal models of the plant, probably by learning inverse dynamics [4, 5]. For many years the consensus has been that the system is still velocity driven. However, there is now evidence that at least the smooth pursuit system receives position as well as velocity commands [6].

It is now firmly established that most oculomotor control is under real-time adaptive control [2,7,8,9]. That is, synaptic weights are constantly being updated to maintain precise oculomotor control. Anatomically, it has been shown that the cerebellum is an important structure that mediates adaptive control. In particular, the flocculus/paraflocculus is important for maintaining ocular plant compensation [2], and it has been proposed that the correct inverse dynamics are learnt from the visual response, which has been assumed to be retinal slip (Fig.1) [5].

A key point for us is that adaptive control implies the existence of a control objective embedded in the brain. It also implies that oculomotor behaviour is under real-time optimal control, where the 'optimum' corresponds to the control objective. Such a system requires signals that code departures from the optimum and teach/guide the system to reach the optimum, and it has long been believed that the climbing fibre input to the cerebellum carries such "teaching signals" [10, 11].

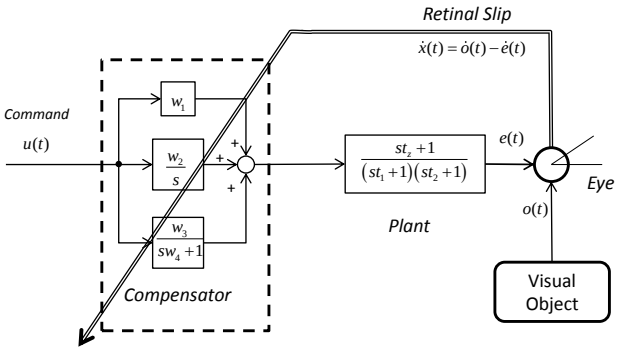

Figure 1: Schematic of the oculomotor final common path. The plant is modeled as a 2nd order linear system with a zero. The motor command $u(t)$, drives a compensator which ideally cancels the poles (t_1, t_2) and zero (t_z). The compensator parameters $w_{1,2,3,4}$ are under adaptive control usually considered to be driven by retinal slip $\dot{x}(t)$ (thick line), which is the difference between object velocity $\dot{o}(t)$ and eye velocity $\dot{e}(t)$.

A. *Pathological Oscillations*

Human eye movement control can be disturbed by neurological disorders (acquired or congenital) and pathological eye movements have been intensely studied [12]. Of particular interest is the existence of pathological oscillatory instabilities of the eyes, called 'nystagmus'. A puzzling question is why nystagmus can be permanent, as one would expect adaptive control to compensate for any instabilities (as indeed can occur in some vestibular disorders). One possibility is that the adaptive controller receives inaccurate information or is itself damaged by disease. This makes sense in some conditions where there is brainstem and cerebellar damage. However, in infantile nystagmus syndrome (INS), infants develop lifelong eye oscillations from birth or soon after. There is no evidence of brain damage, but surprisingly most are born with a visual (sensory) abnormality [13, 14].

There have been many attempts to model INS as a control system [15,16,17,18,19]. Most are variants on the normal velocity control model (fig.1) but fail to take into account that INS is also sensitive to positional error [20]. It appears that, as with the smooth pursuit system, we also need to take into account positional error. Moreover, none of these models explain why adaptive control does not automatically nullify the nystagmus.

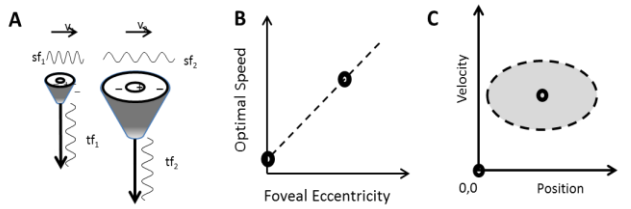

Figure 2: Schematic of visual spatiotemporal tuning. (a) Receptive fields (RFs) are temporally and spatially tuned, and so become tuned to image speed depending on size of RF. (b) Small RFs near the fovea are tuned to low but non-zero speeds. Larger RFs are tuned to higher speeds. (c) It is not physically realizable to maintain optimum speed and position, and some compromise is needed, such as a least-squares (quadratic) cost function (8).

Clearly adaptive control requires visual feedback, and cannot take place in infants until birth. However, vision undergoes considerable post-natal development (especially foveal vision), and it is not clear how this interacts with the adaptive control of eye movements. Harris & Berry [21,22] have proposed that infantile nystagmus persists because, paradoxically, it may fulfill the adaptive control objective (see below), together with anomalous timing of sensory development in infancy [23]. At the heart of this idea is the peculiar arrangement of the visual system which is intrinsically tuned to position and a non-zero image velocity, *p* (quite unlike a conventional camera) (fig.2). Coupled with a mixture of positional and velocity errors, oscillations can emerge as optimal with properties very similar to those observed in INS.

This approach is fundamentally different from the conventional control theory approach since it asks about the adaptive control objective, rather than modeling how oscillations may be driven. However, Harris & Berry [21] only considered the global optimum as a variational problem, without asking whether it could be reached by an admissible motor command. Here we re-frame the problem in terms of optimal control theory to gain more insight into whether or how the plant and its command may constrain adaptive control of the combination of optimizing image position as well as velocity.

II. METHODS

We consider a 1-dimensional problem with horizontal eye position denoted by $e(t)$, with the origin corresponding to the retinal fovea aligned straight ahead (head fixed). We use the dot notation to denote time derivatives, so that eye velocity is denoted by $\dot{e}(t)$ and acceleration by $\ddot{e}(t)$. We denote target position relative to the straight ahead by $o(t)$ and its velocity by $\dot{o}(t)$. Retinal error is denoted by $x(t) = o(t) - e(t)$, and its time derivative (retinal slip) is $\dot{x}(t) = \dot{o}(t) - \dot{e}(t)$.

We assume eye position is controlled by a unity gain linear plant driven by a motor command $u(t)$ (see Fig.1 without a compensator). The plant is assumed to have two real poles with time constants, t_1 and t_2, and a zero with time-constant t_z (see Fig.1), so that:

$$a\ddot{e}(t) + b\dot{e}(t) + e(t) = t_z\dot{u}(t) + u(t) \tag{1}$$

where $a = t_1 t_2$ and $b = t_1 + t_2$. In terms of retinal error:

$$a(\ddot{x} - \ddot{o}) + b(\dot{x} - \dot{o}) + x - o = -t_z\dot{u} - u \tag{2}$$

For a state-space representation, we define the state vector by $\mathbf{x} = [x_1, x_2 \cdots]^T$ where $x_1 \equiv x$, $x_2 \equiv \dot{x}$, etc., which is assumed to be directly observable, and $o_1 \equiv o$, $o_2 \equiv \dot{o}$, etc.. We will use either classical or state representation depending on context.

A state space representation of (2) is $\dot{\mathbf{x}} - \dot{\mathbf{o}} = \mathbf{A}(\mathbf{x} - \mathbf{o}) + \mathbf{B}u$ or explicitly:

$$\begin{bmatrix} \dot{x}_1 - \dot{o}_1 \\ \dot{x}_2 - \dot{o}_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -1/a & -b/a \end{bmatrix} \begin{bmatrix} x_1 - o_1 \\ x_2 - o_2 \end{bmatrix} + \begin{bmatrix} -c/a \\ -d/a \end{bmatrix} u(t) \tag{3}$$

where $c = t_z$, $d = 1 - bt_z/a$. In the simulations below, we set $t_1 = 0.3s$, $t_2 = 0.010s$ $t_z = 0.080s$ based on parameters given by Zee et al. [3]. For a stationary object straight-ahead (i.e. imaged on the foveas when the eye is still), the state-space equation for retinal error is:

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -333.33 & -103.33 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} + \begin{bmatrix} -26.67 \\ -0.00242 \end{bmatrix} u(t). \tag{4}$$

Numerical simulations were carried out with Matlab (Mathworks, USA). The duration of the epoch, *T*, was set to unity. The ideal speed, *p*, was also set to unity. Note that fields of extremals are proportional to *p*.

III. VISUAL LAGRANGIANS AND GLOBAL OPTIMA

Our fundamental assumption is that the control objective of any adaptive control is essentially visual. That is, the

cost/benefit of eye movements is determined by their visual consequences, and reflects the notion that eye movements have no other purpose than to control the orientation of the retinas. We assume that consequences can be written as some scalar which we call 'cost' (i.e. negative benefit).

We therefore introduce the idea of a "visual Lagrangian" which defines how the cost to the visual system depends on the state vector of the image on the retina (and possibly its history) in a particular visual task. We define the cost of an eye movement behaviour over some epoch of time $(0 \le t \le T)$ by a Lagrangian $L(\mathbf{x},t)$ with the expressed assumption that it is independent of control $u(t)$ (and assuming control bounds are not reached). The cost, J, of an eye movement is then given by:

$$J = \int_0^T L(\mathbf{x},t)dt \tag{5}$$

Provided smoothness constraints are met, the optimal trajectory of the retinal image is a solution to the Euler-Lagrange equation (ELE):

$$\frac{\partial L}{\partial x_1} - \frac{d}{dt}\left(\frac{\partial L}{\partial x_2}\right) = 0 \tag{6}$$

The solution to (6) with the minimum cost over the space of admissible functions and subject to boundary conditions (if any) is the optimum denoted by $\mathbf{x}^*(t)$. There is no guarantee that the neuro-muscular system can reach this optimum.

It has usually been assumed (implicitly) that the control objective is to minimise retinal slip, that is, to reduce image speed to zero. This makes sense from an engineering perspective where cameras are the sensory input and their responses are degraded by image motion blur. However, this is a serious oversimplification for biological visual systems. The fundamental (and only) signals emanating from the vertebrate retina are the responses of the retinal ganglion cells (RGCs), whose axons constitute the optic nerve. RGCs are driven by photoreceptors, bipolar, horizontal, and amacrine cells, which filter photon capture over a small optical region in space, called a receptive field (RF). Each RF (and hence RGC) is tuned to be maximally responsive to spatial frequency (similar to a bandpass spatial modulation transfer function), but also tuned to be maximally responsive for particular temporal frequency (similar to a bandpass temporal filter) [23]. This means that each RGC RF is tuned to its own non-zero image speed (see Fig. 2). Thus, each 'pixel' of the biological image is tuned to a position due to its spatial location in the retina, but also to a non-zero image velocity due its RF tuning properties.

In most vertebrates, the diameters of RFs are also not constant across the retina. In humans and other primates, the retina has a "fovea"' with very small and dense RFs, that increase in diameter and decrease in density with radial distance from the fovea (although the overlap of RFs remains roughly constant). Thus the optimal velocity, p, increases with eccentricity from the fovea depending on spatial frequency (Fig.2b) [24].

A. *Pure Motion Lagrangian*

One consequence of this organisation is that retinal stabilization of an image actually reduces contrast (with greater reduction in the peripheral retina). Thus, eliminating retinal slip would be counterproductive. However, Michael Land [25] recognised that the degree of stabilization in animals is frequently better than needed for RFs, and proposed that good stabilization locks the visual background onto the retina to better support detection of an object moving against its visual background (even if the visual contrast of the background is reduced). The control objective of this visual task would be to stabilize the visual background, or large area of the visual scene. Assuming the optimum to be locally quadratic, the simplest Lagrangian would be

$$L = x_2^2 \tag{7}$$

which is obviously minimised with zero retinal slip. Since coherent motion of large areas of the visual scene are usually caused by self-motion, the VOR and OKR seem likely to be the supporting sub-systems.

B. *Visual Contrast Lagrangian*

Prolonged viewing of small or distant stationary objects requires positional control, but perfect stabilization reduces visual contrast. The fundamental problem is that a RGC RF is tuned to non-zero velocity and to a spatial location (by virtue of its position in the retinotopic map) and will have an optimal phase point (see Fig. 2c). The requirement to match ideal motion and position simultaneously cannot be physically realised, as optimal image motion will take the image away from the RF and an optimal image position control would have zero motion.

The image of a visual object will usually stimulate many RFs, where each RF will have its own optimal phase point. This is a complex problem in its own right, but overall we expect there to be some image trajectory that maximises total response of RGCs. Again, assuming quadratic local optimum in position and velocity, the simplest Lagrangian is:

$$L = \alpha x(t)^2 + \beta(\dot{x}(t) - p)^2 \tag{8}$$

where p is the ideal image speed, which is assumed to be constant for small movements. Here, α and β are constants that weight position and velocity error. From (6), the optimal state trajectory is then:

$$x^*(t) = A\exp(-t/\tau) + B\exp(t/\tau), \qquad \tau = \sqrt{\beta/\alpha} \tag{9}$$

where A and B are undetermined constants. Without any boundary conditions, the *global* optimum (identified by double asterisk) is given by [21]:

$$x^{**}(t) = A^{**}\exp(-t/\tau) + B^{**}\exp(t/\tau)$$

$$A^{**} = \frac{p\alpha\exp(T/\tau)}{1+\exp(T/\tau)} \tag{10}$$

$$B^{**} = \frac{p\alpha}{1+\exp(T/\tau)}$$

The global optimum is illustrated in fig.3a (no zero) and fig.3c for τ=0.45. Thus, an ideal adaptive controller with the capacity to find the global optimum should always generate image motion depending on *p*. This is a direct result of the velocity tuning of RGCs. Even for normal adults, foveal RGCs prefer a small but non-zero image motion, and so-called "steady fixation" are observed to consist of slow alternating drifts and microsaccades [26]. We argue that in some infants with delayed visual development, the oscillations become extreme, which we call "nystagmus" [23].

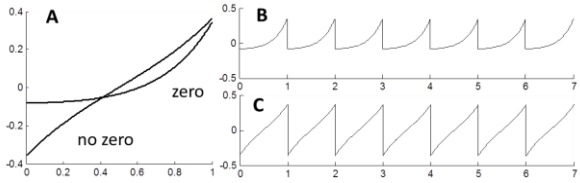

Figure 3: a) Global optima for uncompensated zero ('zero') and compensated zero ('no zero'). b) Ideal position waveform with zero and instantaneous reset; c) same as (b) but for no zero. ***T*=1; *p*=1; τ=0.45**

IV. ADAPTIVE CONTROL & LOCAL MINIMA

We now consider an adaptive process that attempts to minimise the visual Lagrangian. However this process is implemented (e.g. filter, transform, feedback), it must operate on the control signal, $u(t)$, in such a way as to optimise the output objective. Such a process must be constrained by the dynamics of the plant, and so we apply standard optimal control theory (OCT) [27].

We construct a Hamiltonian from the Lagrangian and plant dynamics (without any compensator):

$$H(\mathbf{x},u,t) = L(\mathbf{x},t) + \mathbf{\Gamma}(t)^T[A\mathbf{x} + Bu] \tag{11}$$

where $\mathbf{\Gamma}(t) = \left[\lambda_1(t),\cdots,\lambda_n(t)\right]$ is the co-state vector. Because $H(\mathbf{x},u,t)$ is a linear function of control, the optimal strategy must be singular or bounded (bang-bang) control. We assume that the motor command is not bounded over the range of eye movements considered here. For a 2nd order system with a zero, the expanded Hamiltonian is:

$H(\mathbf{x},u,t) = L(\mathbf{x},t) + \lambda_1(x_2 + cu/a) + \lambda_2(-x_1/a - bx_2/a + du/a)$, and from OCT, the optimum must satisfy:

$$\begin{aligned}
\frac{\partial H}{\partial x_1} &= -\dot{\lambda}_1 = \frac{\partial L}{\partial x_1} - \lambda_2/a \\
\frac{\partial H}{\partial x_2} &= -\dot{\lambda}_2 = \frac{\partial L}{\partial x_2} + \lambda_1 - b\lambda_2/a \\
\frac{\partial H}{\partial u} &= 0 = c\lambda_1/a + d\lambda_2/a.
\end{aligned} \tag{12}$$

Eliminating $\dot{\lambda}_1$ and $\dot{\lambda}_2$:

$$\lambda_1 = \frac{-1}{q}\left(\frac{c}{d}\frac{\partial L}{\partial x_1} + \frac{\partial L}{\partial x_2}\right), \qquad q = 1 + \frac{cb}{ad} + \frac{c^2}{ad^2} \tag{13}$$

which yields the differential equation

$$\frac{c}{qd}\frac{d}{dt}\left(\frac{\partial L}{\partial x_1}\right) + \frac{1}{q}\frac{d}{dt}\left(\frac{\partial L}{\partial x_2}\right) + \left(\frac{c^2}{aqd^2} - 1\right)\frac{\partial L}{\partial x_1} + \frac{c}{adq}\frac{\partial L}{\partial x_2} = 0 \tag{14}$$

In general, a function that satisfies this equation does not satisfy the global optimum ELE (6), except for degenerate cases. However, when there is no zero (c=0) (14) is equivalent to (6) and the global optimum can be reached. In other words, provided the plant is linear and there is no zero, the global optimum can be reached by a suitable control signal $u(t)$. If the zero is uncompensated (or the plant is not linear), the global optimum cannot be reached generally (excepting degenerate cases).

As an example, let us return to Harris & Berry's visual Lagrangian with an uncompensated zero. For $L = \alpha x_1^2 + \beta\left(x_2 - p\right)^2$ we have:

$$\frac{\tau^2}{q}\frac{d^2x}{dt^2} + \left(\frac{c}{qd} + \frac{c\tau^2}{adq}\right)\frac{dx}{dt} + \left(\frac{c^2}{aqd^2} - 1\right)x = \frac{\tau^2 pc}{adq} \tag{15}$$

which has the general solution:

$$x(t) = A\exp(r_1 t) + B\exp(r_2 t) + D \tag{16}$$

where r_1, r_2 are the two roots, *A*, *B* are undetermined constants, and $D = \frac{\tau^2 pcd}{c^2 aqd^2 - 1}$. The global optimum for a zero (t_z = 0.08; see sect II) is shown in Fig. 3a and b ('zero'). As can be seen, the optimal trajectory becomes much more curvilinear with a non-compensated zero.

A. Bidirectional Cost

So far, we have assumed unidirectional sensitivity to the optimal speed, *p*. However, it seems reasonable to assume that contrast is independent of the sign of *p*, and that the RF responds equally to leftward and rightward motion. To solve this problem, we consider the function $P(t)$, such that $|P(t)| = p$ but can flip sign *n* times during the epoch, denoted by t_i, $i = 1,n$ where *n* is finite. The quadratic Lagrangian becomes $L = \alpha x(t)^2 + \beta\left(\dot{x}(t) - P(t)\right)^2$. Now, during each of the periods between flipping:

$\left(0 \le t < t_1\right)$, $\left(t_1 < t < t_2\right),\cdots\left(t_{n-1} < t < t_n\right),\left(t_n < t \le T\right)$, $P(t)$ is a constant, and $dP/dt = 0$. At $t = t_i$, the Lagrangian is indeterminate but finite. In the infinitesimal limit, it must follow that the area under the Lagrangian at t_i approaches zero, so that $\sum_{i=1}^{n} L(x(t_i),t_i) = 0$. The cost is therefore given by

$$J = \int_0^T L(x,t), \quad t \ne t_i. \tag{17}$$

We now have

$$\frac{\tau^2}{q}\frac{d^2x}{dt^2} + \left(\frac{c}{qd} + \frac{c\tau^2}{adq}\right)\frac{dx}{dt} + \left(\frac{c^2}{aqd^2} - 1\right)x = \frac{\tau^2 P(t)cd}{c^2 aqd^2 - 1} \tag{18}$$

which is 2nd order as before, but now *p* is replaced by $P(t)$:

$$x(t) = A\exp(r_1 t) + B\exp(r_2 t) + \frac{\tau^2 P(t)cd}{c^2 aqd^2 - 1} \tag{19}$$

Defining $x^+ = x(t)$ when $P(t) = +p$, and $x^- = x(t)$ when $P(t) = -p$, then the constants A and B are chosen together with any boundary conditions, so that

$$J = \int_0^T \min\left[L(x^+(t),t), L(x^-(t),t)\right]dt \qquad (20)$$

which can be solved by a simple search algorithm.

V. BOUNDARY CONDITIONS

The global optima, with or without a plant zero, result in a net change in position. Thus, for a finite *T*, it is necessary to reset position to repeat a cycle. Intuitively, a step reset would seem the optimal way to reset, but this cannot be achieved without infinite control. The fastest reset available to the oculomotor system is the saccade. Saccades take finite time and have end-point positional error, which increases with amplitude. The triggering of saccades is also extremely variable with highly skewed latency distributions, even during a repetitive behaviour such as optokinetic nystagmus [28]. Thus, saccades come with cost that increases with the size of the reset. As pointed out by Harris & Berry [21], the initial uncertainty of a cycle will affect the extremal, which in turn will affect the next reset, possibly leading to 1st order Markov sequences. Given that saccade amplitude is also under adaptive control, the problem is difficult to unravel.

Here we consider various reset procedures, which are equivalent to placing costs on boundary conditions [29, 30], including the absence of a reset, which we call type "0", pure position change: type "1", and position and velocity change: type "2". Initially, we consider a compensated zero.

Type 0

One possibility is that the cost of making a saccade is so high that it pays not make the reset at all, so that $x_1(0) = x_1(T)$. Since $x_1(0) = A + B + D$ and $x_1(T) = A\exp(r_1T) + B\exp(r_2T) + D$, this requires $B = -A\left(1-\exp(r_1T)\right)/\left(1-\exp(r_2T)\right)$, and leaves one degree of freedom. We computed the field of extremals (FOE) by finding the local optimal trajectory for different initial positions, where $A = [x_1(0) - D]/[1-\left(1-\exp(r_1t)\right)/\left(1-\exp(r_2t)\right)]$ based on bidirectional cost function (otherwise no movement would be optimal). As seen in fig.4a, the optimum requires a symmetrical curvilinear positional trajectory. This trajectory also requires an abrupt change in velocity at the end of the cycle (fig.4b). This is strictly not achievable for a 2nd order plant as it would require infinite acceleration. A finite acceleration pulse would 'round-off' the velocity change at the end of the cycle, but whether the saccadic mechanism can accomplish this is not known (and not previously asked).

Type 1

In most eye movement studies, saccades are required to change eye position but not eye velocity. That is, the velocity pulse begins and ends at the same value. This requires $x_2(0) = x_2(T)$ so that $B = -A\left(1-r_1\exp(r_1T)\right)/\left(1-r_2\exp(r_2T)\right)$ and therefore $A = [x_1(0) - D]/[1-\left(1-r_1\exp(r_1t)\right)/\left(1-r_2\exp(r_2t)\right)]$ The FOE for this scenario shows curvilinear trajectories and includes the global optimum (for no zero) (fig.4c). This requires discontinuity in acceleration at the end of the cycle (fig.4d), but this is feasible with a 2nd order plant.

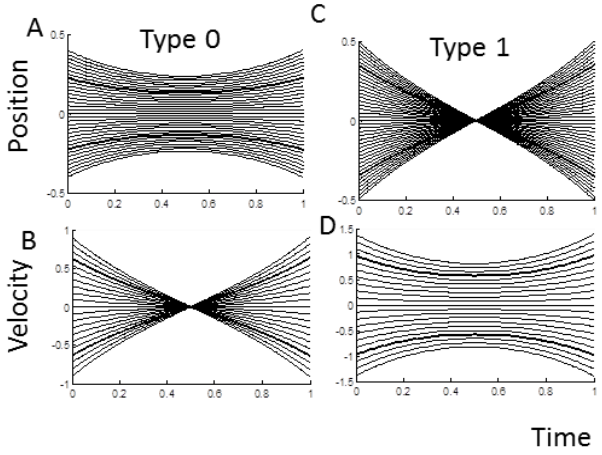

Figure 4: Fields of extremals for type 0 and type 1 boundary conditions with bi-directional cost. Thick lines are local optima. ***T*=1; *p*=1; *τ*=0.45**

Type 2

The final possibility is that saccades can change both velocity and position simultaneously (that is, they are true state changers). For a specified $x_1(0)$, we have $B = x_1(0) - A$ (note *D*=0 for no zero). This FOE (unidirectional cost) includes the global optimum (fig.5a), but for other extremals, a change in velocity at the end of the cycle is needed.

These FOE's also depend on the zero. If the zero is not compensated, the type 0 and 1 FOE's remain qualitatively similar (not shown). However, for type 2, an uncompensated zero leads to a quite different field with extremals starting very slowly before accelerating later (figs.5c,d). Presumably this occurs because of the reduced constraints in type 2.

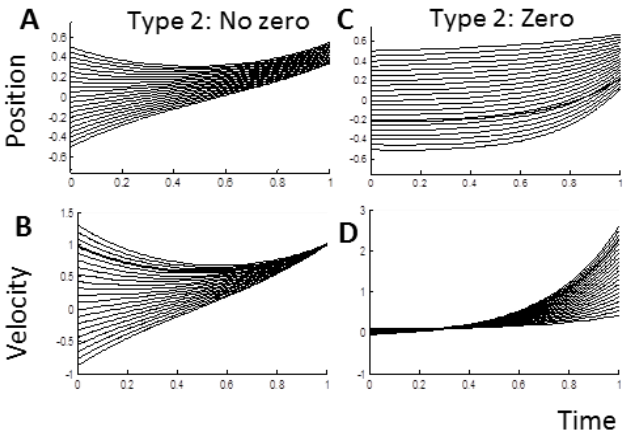

Figure 5: Fields of extremals for type 2 boundary conditions with unidirectional cost and with zero compensated (left) and uncompensated (right). ***T*=1; *p*=1; *τ*=0.45**

VI. DISCUSSION

We propose that OCT can provide explanatory insight into the study of eye movements by placing bounds on how

biological control could evolve and operate. We have made three fundamental assumptions about eye movement behaviour.

First, control is adaptive in real-time. There is ample evidence for this, but it implies that there is some 'goal' or control objective that is being constantly pursued by the adaptive controller. How this is encoded in the brain is unknown, but it is tempting to argue that RGC firing patterns across the retina code a 'raw' field (including motion vectors) of visual benefit and cost. This raw field is refined by subcortical and cortical processing to provide teaching signals to the climbing fibre input to the cerebellum, which then finds the optimal motor control (adaptive control).

Second, we have assumed an ideal adaptive controller, so that observed behaviour meets the objective. Of course, as in man-made systems, adaptive controllers are noisy and can fail if they are damaged or receive incorrect information. However, in normal healthy brains, this assumption makes sense because we would expect evolution (natural selection) to come as close as possible to an ideal controller.

Third, we argue that the control objective is purely visual (sensory), and not motor. This may seem unusual in design engineering, where costs are often attached to the motor control (fuel, energy). However, eye movements only serve the function of controlling the orientation of the retinas in time. This leads to the idea of visual Lagrangians, $L(\mathbf{x},t)$, which are functions of the state vector and time, but not control. So far, we have only considered simple visual Lagrangians. In reality, we expect more complex functionals of state, as eye movements have higher dimensions (horizontal, vertical and torsional for each eye). The visual impulse response function should impose an explicit dependence on absolute time relative to the last saccade (or stimulus onset). The visual characteristics (size, shape, contrast) of the object's image, and its eccentricity relative to the fovea need to be taken into account. Nevertheless, the optimal state vector represents the global optimum $\mathbf{x}^{**}(t)$, and can be found either analytically (variational calculus) or numerically.

Of course, the brain can only ever find the optimum by manipulating the motor control input $u(t)$, which affects the state vector via the motor plant. The goal of the adaptive controller is to operate on $u(t)$ (via some controller/ compensator) so that the objective can be reached. We believe the fundamental problem confronting biological motor systems is that the dynamics of the plant can prevent the global optimum from being reached. The power of OCT is that it tells us what the best outcome could be. The Hamiltonian is the key to this problem since the optimal control is given by minimising $H(\mathbf{x},u,t)$ with respect to $u(t)$. The Hamiltonian always depends on $u(t)$, even though the visual Lagrangian does not. Clearly, from (5), for the optimal control to reach the global optimum, the plant needs to be linear (over the range of control needed to find the optimum), and not bounded. Thus, singular control is optimal. However, less obvious is that, even for a linear plant, a zero in the plant prevents the global optimum from being reached (apart from degenerate cases).

Eye movements are typically a sequence of slow eye movements separated by saccades. Saccades are, therefore, effectively transitions between singular arcs. Saccade amplitude and the transfer of state from a saccade to the subsequent slow eye movement are known to be under adaptive control [2,9]. Although finding the optimal transition between singular arcs is a well-studied non-trivial problem [27], how this can be achieved with saccades remains an important and fundamental theoretical issue. Saccades cannot be under singular control because of the presence of signal-dependent motor noise [31,32]. Due to this noise and also targeting error, saccade accuracy decreases with amplitude and will lead to uncertainty in the starting point of the next singular arc and forces a variable boundary condition. The ideal adaptive controller would take the start position into account and follow the appropriate optimal trajectory (i.e. encode FOEs). However, to what extent the saccade trajectory can be manipulated is poorly understood. Conventionally, saccades are thought to change position but not velocity, which would constrain FOEs (Fig.4c). If velocity can also be changed, then different boundary conditions can be exploited.

INS is a unique developmental disorder of eye movements. We have proposed that the oscillations are consistent with an adaptive controller with an objective to minimise the quadratic visual Lagrangian (8) [21]. Here we argue that this is essentially the same as the eye drifts and microsaccades seen in normal individuals, but for a larger optimal image speed, *p*. Slow drifts and nystagmus slow phases may be driven by different mechanisms, but they seem to share the same control objective. Interestingly, the ability to reach the global optimum requires the zero to be compensated, and it is tempting to speculate that there is a sequence of adaptive controls. Initially, the zero is compensated via a simple Lagrangian such as minimising image velocity (retinal slip), as has been proposed for the vestibular system [5]. Once achieved, a second adaptive system minimises the Lagrangian with positional error, possibly via the smooth pursuit system. This sequence is consistent with known development, and may provide an example of steering control around local minima, as recently proposed by Harris [23].

Multiple adaptive controllers lead to considerable complexity, which would be extremely difficult to model conventionally. However, OCT allows us to explore their effects, even without knowing their details. We have only examined a rather simple system, but the possibility that we may be able to elucidate sequences of adaptive phases is exciting, as it may provide a handle on understanding the underlying principles of biological development and learning.

VII. CONCLUSIONS

From the 1960's, the principles of control systems theory have provided insights into how Nature might control eye movements, which have been confirmed time and time again by empirical physiological and behavioural experiments. One is led to the conclusion that Nature must be an 'excellent' engineer. Perhaps we should not be surprised.

Good engineering designs are based on well-tried, robust and optimal performance – exactly the qualities we expect from systems that have evolved under natural selection. Our view is that the principles of optimal control theory provide insights into why oculomotor systems have evolved the way they are, and in particular, provide a framework for understanding development and adaptive control.

Optimising vision (via visual Lagrangians) generally requires singular optimal control, which may be unique in biological motor control. This requires a linearized plant with compensation for zeroes, which is consistent with empirical observation. How this is learnt has not been previously asked, but may be highly relevant to normal and pathological development. As saccadic eye movements are under non-singular control, it is particularly important both empirically and theoretically to understand how transitions between singular and non-singular arcs are accomplished. We are currently investigating this non-trivial question.

references

D.A. Robinson, "The mechanics of human saccadic eye movement," J. Physiol. Lond, vol. 174, no. 2, pp. 245-264, 1964

L.M. Optican, and F.A. Miles, "Visually induced adaptive changes in primate saccadic oculomotor control signals," J. Neurophysiol, vol. 54, no. 4, pp. 940-958, 1985.

D.S. Zee, E.J. Fitzgibbon, and L.M. Optican, "Saccade-vergence interactions in humans," J. Neurophysiol, vol. 68, no. 5, pp. 1624-1641, 1992.

A.M. Green, H. Meng, and D.E. Angelaki, "A reevaluation of the inverse dynamic model for eye movements," J. Neurosci, vol. 27, no. 6, pp. 1346-1355, 2007.

J. Porrill, and P. Dean, "Cerebellar motor learning: when is cortical plasticity not enough?" PLoS. Comput. Biol, vol. 3, no. 10, pp. 1935-1950, 2007.

G. Blohm, M. Missal, and P. Lefevre, "Direct evidence for a position input to the smooth pursuit system," J. Neurophysiol, vol. 94, no. 1, pp. 712-721, 2005.

A. Gonshor, and G.M. Jones, "Short-term adaptive changes in the human vestibule-ocular reflex arc," J. Physiol, vol. 256, no. 2, pp. 361-379, 1976.

A. Gonshor, and G.M. Jones, "Extreme vestibule-ocular adaptation induced by prolonged optical reversal of vision," J. Physiol, vol. 256, no. 2, pp. 381-414, 1976.

F.R. Robinson, R. Soetedio, and C. Noto, "Distinct short-term and long-term adaptation to reduce saccade size in monkey," J. Neurophysiol, vol. 96, no. 3, pp. 1030-1041, 2006.

D. Marr, "A theory of cerebellar cortex," J. Physiol. Lond, vol. 202, no. 2, pp. 437-470, 1969.

J.S. Albus, "A theory of cerebellar function," Math. Biosci, vol. 10, no. 1-2, pp. 25-61, 1971.

R.J. Leigh, and D.S. Zee, The Neurology of Eye Movements, 4th ed., Oxford, Oxford University Press, 2006.

I. Casteels, C.M. Harris, F. Shawkat, and D. Taylor, "Nystagmus in infancy," Br. J. Ophthalmol, vol. 76, no. 7, pp. 434-437, 1992.

B. Lorenz, and E Gampe, "Analyse von 180 patienten mit sensorischem Defektnystagmus (SDN) und kongenitalem idiopathischem Nystagmus (CIN)," Klin. Monbl. Augenheilkd, vol. 218, no. 1, pp. 3-12, 2001.

D.S. Zee, R.J. Leigh, and F. Mathieu-Millaire, "Cerebellar control of ocular gaze stability," Ann. Neurol, vol. 7, no. 1, pp. 37-40, 1980.

L.M. Optican, and D.S. Zee, "A hypothetical explanation of congenital nystagmus," Biol. Cybern, vol. 50, no. 2, pp. 119-134, 1984.

R.J. Tusa, D.S. Zee, T.C. Hain, and H.J. Simonsz, "Voluntary control of congenital nystagmus," Clin. Vis. Sci, vol. 7, no. 3, pp. 195-210, 1992.

C.M. Harris, "Problems in modelling congenital nystagmus: towards a new model," in Eye Movement Research: Mechanisms, Processes, and Applications, pp.239-253, Amsterdam, Elsevier, 1995.

D.S. Broomhead, R.A. Clement, M.R. Muldoon, J.P. Whittle, C. Scallan, R.V. Abadi, "Modelling of congenital nystagmus waveforms produced by saccadic system abnormalities," Biol. Cybern, vol. 82, no. 5, pp. 391-399, 2000.

G Kommerell, "Congenital nystagmus: control of slow tracking movements by target offset from the fovea," Graefes. Arch. Clin. Exp. Ophthalmol, vol. 224, no. 3, pp. 295-298, 1986.

C.M. Harris, and D.L. Berry, "A distal model of congenital nystagmus as nonlinear adaptive oscillations," Nonlinear. Dynam, vol. 44, no. 1-4, pp. 367-380, 2006.

C.M. Harris, and D. Berry, "A developmental model of infantile nystagmus," Semin. Ophthalmol, vol. 21, no. 2, pp. 63-69, 2006.

C.M. Harris, "Oculomotor developmental pathology: an evo-devo perspective," in Oxford Handbook of Eye Movements, pp. 663-686, Oxford, Oxford University Press, 2011.

V. Virsu, J. Rovamo, P. Laurinen, and R. Nasanen, "Temporal contrast sensitivity and cortical magnification," Vis. Res, vol. 22, no. 9, pp. 1211-1217, 1982.

M.F. Land, and D-E. Nilsson, Animal Eyes, 1st ed., Oxford: Oxford University Press, 2002.

M. Rolfs, "Microsaccades: small steps on a long way," Vis. Res, vol. 49, no. 20, pp. 2415-2441, 2009.

A.E. Bryson, and Y.C. Ho, Applied Optimal Control, Washington DC, Hemisphere, 1975.

J. Waddington and C.M. Harris, "The distribution of quick phase interval durations in human optokinetic nystagmus," Exp. Brain Res, vol. 224, no. 2, pp.179-187, 2013.

C.M. Harris, "Exploring smoothness and discontinuities in human motor behaviour with Fourier analysis," Math. Biosci, vol. 188, no. 1-2, pp. 99-116, 2004.

C.M. Harris, and M.R. Harwood, "Boundary conditions in human movement III: Fourier analysis of reaching. II: The Fourier approach. I: Constraining constraints," in Proceedings of the IASTED International Conference Biomedical Engineering, pp. 629-640, Anaheim, ACTA Press, 2005.

C.M. Harris, and D.M. Wolpert, "Signal-dependent noise determines motor planning," Nature, vol. 394, no. 6695, pp. 780-784, 1998.

C.M. Harris, and D.M. Wolpert, "The main sequence of saccades optimizes speed-accuracy trade-off," Biol. Cybern, vol. 95, no. 1, pp. 21-29, 2006.

Chris Harris obtained a BSc (Hons) in Physics in 1973, and an MSc in Optics 1974 from Imperial College, London. After working as an optoelectronics engineer, he emigrated to New York where he obtained his PhD in 1986 in the study of human eye movement development. In 1989 he returned to the UK and set up an eye movement laboratory in Great Ormond Street hospital carrying out clinical research into abnormal oculomotor

development. In 2000, he was appointed Professor of Neuroscience at the Plymouth University where he continues experimental and theoretical research into visual and motor development in natural and artificial systems. He has published over 100 papers in the field.

Jonathan Waddington was born in Manchester on July 19 1979. He graduated with honours in Physiology at UCL in 2001, and completed his PhD in Neuroscience at Plymouth University in 2012 with a thesis on optokinetic nystagmus. He is currently a post-doctoral research associate working on a knowledge transfer partnership between the University of Lincoln and the WESC Foundation to implement a computer game to improve sight for visually impaired children. He is an author of papers on the application of distribution analysis and linear stochastic models to eye movement data.