Optimal Feedback Control of Gaze

Final Project for
580.691 - Learning Theory

Azwad Sabik

May 9, 2016

Introduction

Direction of the human gaze is an endeavor that involves coordination between the motion of both the eyes and the head. When we direct our gaze upon a particular target in our field of view, the behavior begins with a motion of our eyes (a saccade), is followed up with a slower motion of our head in the same direction, and a concludes with the eyes returning to their resting state at the center of our head while as it continues to orient itself towards the target. This natural behavior implies that the gaze is executed not just with a goal of reaching the target - it is important that the eccentricity of the eyes is minimized as well. In addition, perturbation of this behavior through head immobilization results in a change in saccade magnitude and duration, implying feedback between the head and eye motions dictating control of the overall system. It is in light of these observations that we attempt to simulate the motion and control of a gaze through the use of optimal feedback control in this project.

We utilize a simple one-dimensional model for the head and eye system, treating both eye and head positions as one-dimensional rotations about the body's longitudinal axis, and set the resulting gaze position equal to their sum. The position of the target in our field of view is modeled similarly as a one-dimensional rotation. The dynamics of both the eyes and head are modeled as spring-mass-damper systems evolving over discrete time-steps, with their motions determined by the torque applied to each.

In order to simulate evolution of the system under optimal feedback control, we calculate optimal Kalman and feedback gains corresponding to a fixed movement duration and a desired time for arrival of the gaze at the target position. Through these, we can simulate both the system's estimation of its own pose (position, velocity, and applied torque for both the head and the eyes) as well as the motor commands (muscle activations) it might provide to each given a particular set of costs associated with the action and state at each time step. Specifically, we associate costs with deviation of the the gaze position from the goal position, deviation of the eye position from its resting position, and effort associated with motor inputs. Pose estimates and motor inputs are considered optimal with feedback when the gains which generate them will attempt to minimize these costs regardless of the state in which they are utilized, as per the basis of the Kalman filter and the Bellman principle of optimality.

The goal of this project is to demonstrate the effectiveness of utilizing this type of optimization in simulating the aforementioned natural behavior of a gaze, and additionally emphasize the importance of accounting for both signal-dependent noise and feedback while optimizing. It is to this latter end that simulations are carried out with five sets of parameters, including two which serve to demonstrate the effects of accounting for signal-dependent noise, and two which demonstrate the effects of perturbing the behavior by immobilizing the head for part of the simulation.

Methods

This section describes the functions generated as components of the simulation and corresponding experiments. Each of the components was written and tested with MATLAB R2015b. See the project attachments for source.

calculateAandB.m

The dynamics of our head and eye models both are described by the following continuous-time system:

$$\begin{bmatrix} \dot{x_1} \\ \dot{x_2} \\ \dot{x_3} \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ -\frac{k}{m} & -\frac{b}{m} & \frac{1}{m} \\ 0 & 0 & -\frac{\alpha_2}{\alpha_1} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \frac{1}{\alpha_1} \end{bmatrix} u$$

In these models, x_1 is the position, x_2 is the velocity, and x_3 is the applied torque. The head and eye models vary only by their values for m, b, and α_1. Experimentally measured time constants of each system are used to compute the differing system parameters, with $b = \tau_1 + \tau_2$, $m = \tau_1\tau_2$, $\alpha_1 = \tau_3$. The remaining parameters are held constant at $k = \alpha_2 = 1$. These systems can be more simply represented as:

$$\dot{x}_e = A_e x_e + b_e u_e$$
$$\dot{x}_h = A_h x_h + b_h u_h$$

for the eyes and head respectively. We can stack the coefficient matrices in each of these systems to produce a single, combined continuous-time system:

$$A_c = \begin{bmatrix} A_e & 0 \\ 0 & A_h \end{bmatrix}$$

$$B_c = \begin{bmatrix} b_e & 0 \\ 0 & b_h \end{bmatrix}$$

and reformulate this system as one that operates in distinct time-step Δ:

$$A = \begin{bmatrix} exp(A_c\Delta) & 0_{6\times1} \\ 0 & 1 \end{bmatrix}$$

$$B_c = \begin{bmatrix} A_c^{-1}(exp(A_c\Delta) - I)B_c \\ 0_{1\times2} \end{bmatrix}$$

The overall system can then be represented as

$$x^{(k+1)} = Ax^{(k)} + Bu^{(k)} + \varepsilon_x^{(k)} + B\sum_i C_i u^{(k)}\phi_i^{(k)}$$
$$y^{(k)} = Hx^{(k)} + \varepsilon_y^{(k)}$$

where

$$H \equiv \begin{bmatrix} -1 & 0 & 0 & -1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

which implies the sensed values in the system are the position of the eye in the head and position of the target relative to the gaze. C_i is a parameter corresponding to a noise dependence on the input signal u, and

$$\varepsilon_x \in N(0, Q_x)$$
$$\varepsilon_y \in N(0, Q_y)$$
$$\phi \in N(0, 1)$$
$$\mu \in N(0, 1)$$

calculateKalmanGains.m

In modeling a feedback system that does not 'know' to its true state, it is necessary to generate estimates of the state using information from whatever sensor feedback is available - in this case of the head-eyes system, it is assumed that only the position of the eye in the head and the position of the target relative to the gaze are known. The Kalman filter can be used to estimate the state given these pieces of sensor information, but in systems affected by signal- dependent noise, the Kalman gains are dependent upon the system's state estimates. Thus, in order to plan Kalman gains for the duration of a movement, it is necessary to predict the state estimates throughout - and in the case where state transitions are the result of feedback gains, the computations become dependent on an initial prediction of those gains. We later see that computation of optimal feedback gains depends upon the Kalman gains for the planned motion, and that the computation of these gains can be carried out iteratively, with feedback gains $G^{(k)}$ initialized to zero matrices. Given this, Kalman gains $K^{(k)}$ are computed as follows:

$$K^{(k)} = S_e^{(k|k-1)} H^T (H S_e^{k|k-1} H^T + Q_y)^{-1}$$

$$\begin{aligned} S_e^{(k|k+1)} = {} & A(I - K^{(k)} H) S_e^{(k|k-1)} A^T + Q_x + \cdots \\ & \sum_i B C_i G^{(k)} S_x^{(k|k-1)} G^{(k)^T} C_i^T B^T \end{aligned}$$

$$\begin{aligned} S_x^{(k+1|k)} = {} & (A - BG^{(k)}) S_x^{(k|k-1)} (A - BG^k)^T + \cdots \\ & (A - BG^{(k)}) S_{xe}^{(k|k-1)} (AK^{(k)} H^T) + \cdots \\ & AK^{(k)} H S_{xe}^{(k|k-1)^T} (A - BG^{(k)})^T + \cdots \\ & AK^{(k)} H S_e^{(k|k-1)} A^T \end{aligned}$$

$$S_{xe}^{(k+1|k} = (A - BG^k) S_{xe}^{(k|k-1)} (I - K^{(k)} H)^T A^T$$

calculateFeedbackGains.m

Optimal feedback gains essentially provide the system with a means of selecting a cost-minimizing input signal at each time-step given the current state estimate:

$$u^{(k)} = -G^{(k)} \hat{x}^{(k|k-1)}$$

These gains are computed through repeated application of the Bellman equation, with the cost per step $a^{(k)}$ modeled as

$$a^{(k)} = x^{(k)T}T^{(k)}x^{(k)} + u^{(k)T}L^{(k)}u^{(k)} + \frac{\lambda\beta}{1+\lambda\beta}$$

where T and L represent accuracy and effort costs respectively. This can be reformulated to produce a value function of the following form:

$$v_{\pi^*}(x^{(k)}, \hat{x}^{(k)} = x^{(k)^T}W_x^{(k)}x^{(k)} + (x^{(k)} - \hat{x}^{(k)})^T W_e^{(k)}(x^{(k)} - \hat{x}^{(k)}) + w^{(k)}$$

ultimately allowing for the following computations:

$$G^{(k)} = (L + C_x^{(k+1)} + C_e^{k+1} + B^T W_x^{(k+1)} B)^{-1} B^T W_x^{(k+1)} A$$

$$W_e^k = (A - AK^k H)^T W_e^{k+1}(A - AK^{(k)}H) + G^{(k)^T} B^T W_x^{(k+1)} A$$

$$W_x^k = H^T T^{(k)} H + A^T W_x^{(k+1)} A - G^{(k)^T} B^T W_x^{(k+1)} A$$

$$C_x^{(k+1)} = \sum_i C_i^T B^T W_x^{(k+1)} B C_i$$

$$C_e^{(k+1)} = \sum_i C_i^T B^T W_e^{(k+1)} B C_i$$

The gains should be computed iteratively starting at the last time step, assuming that gain is optimally a zero matrix at that step.

simulateSystem.m

Once Kalman gains and feedback gains have been computed iteratively, it is possible to simulate the evolution of the system. Evolution of the system's state and sensor information can be updated at each time-step using

$$x^{(k+1)} = Ax^{(k)} + Bu^{(k)} + \varepsilon_x^{(k)} + B\sum_i C_i u^{(k)} \phi_i^{(k)}$$
$$y^{(k)} = Hx^{(k)} + \varepsilon_y^{(k)}$$

given an optimal input computed via

$$u^{(k)} = -G^{(k)}\hat{x}^{(k|k-1)}$$

Each time-step's prior state estimate can be found using

$$\hat{x}^{(k+1|k)} = Ax^{(k|k)} + Bu^{(k)}$$

and given sensor information, the state estimate can be updated to the posterior estimate using

$$\hat{x}^{(k|k)} = \hat{x}^{(k|k-1)} + K^{(k)}(y^{(k)} - H\hat{x}^{(k|k-1)})$$

simulateCase.m

Given all of the aforementioned functions, it is possible to simulate and observe the time-evolution under various conditions - in this project,five cases are investigated:

- CASE 1: no signal-dependent noise
- CASE 2: small signal-dependent noise
- CASE 3: large signal-dependent noise
- CASE 4: small signal-dependent noise, head held briefly
- CASE 5: small signal-dependent noise, head held longer

This function is utilized to generate state outputs for a single simulation and a 50-simulation average for a particular case. It achieves this by generating optimal Kalman and feedback gains corresponding to the case by iterating through their computation ten times each, and then using those stable gains as inputs to simulateSystem.m. The simulation is run fifty times and visualized as seen in the following section.

Results

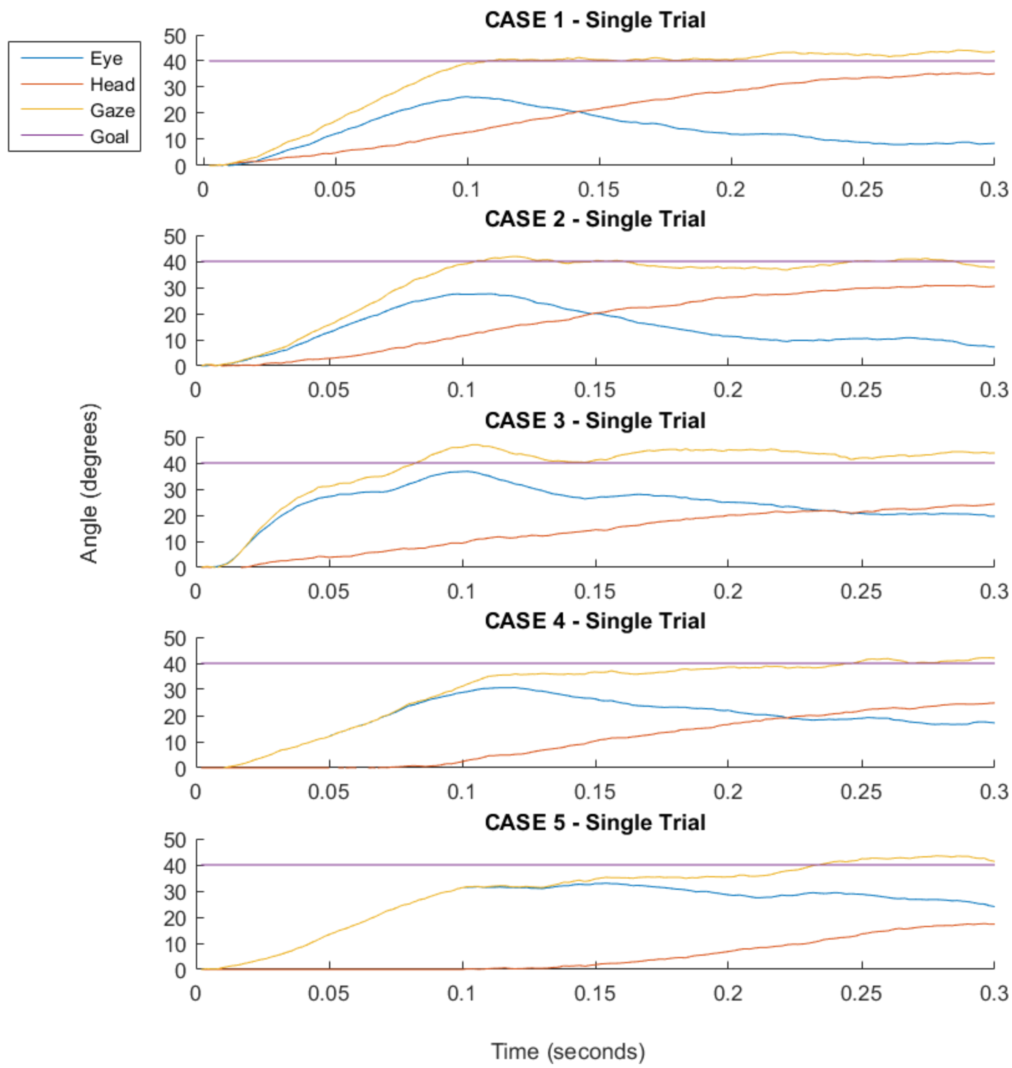

CASE 1 - Single Trial
CASE 2 - Single Trial
CASE 3 - Single Trial
CASE 4 - Single Trial
CASE 5 - Single Trial
Eye
Head
Gaze
Goal
Angle (degrees)
Time (seconds)

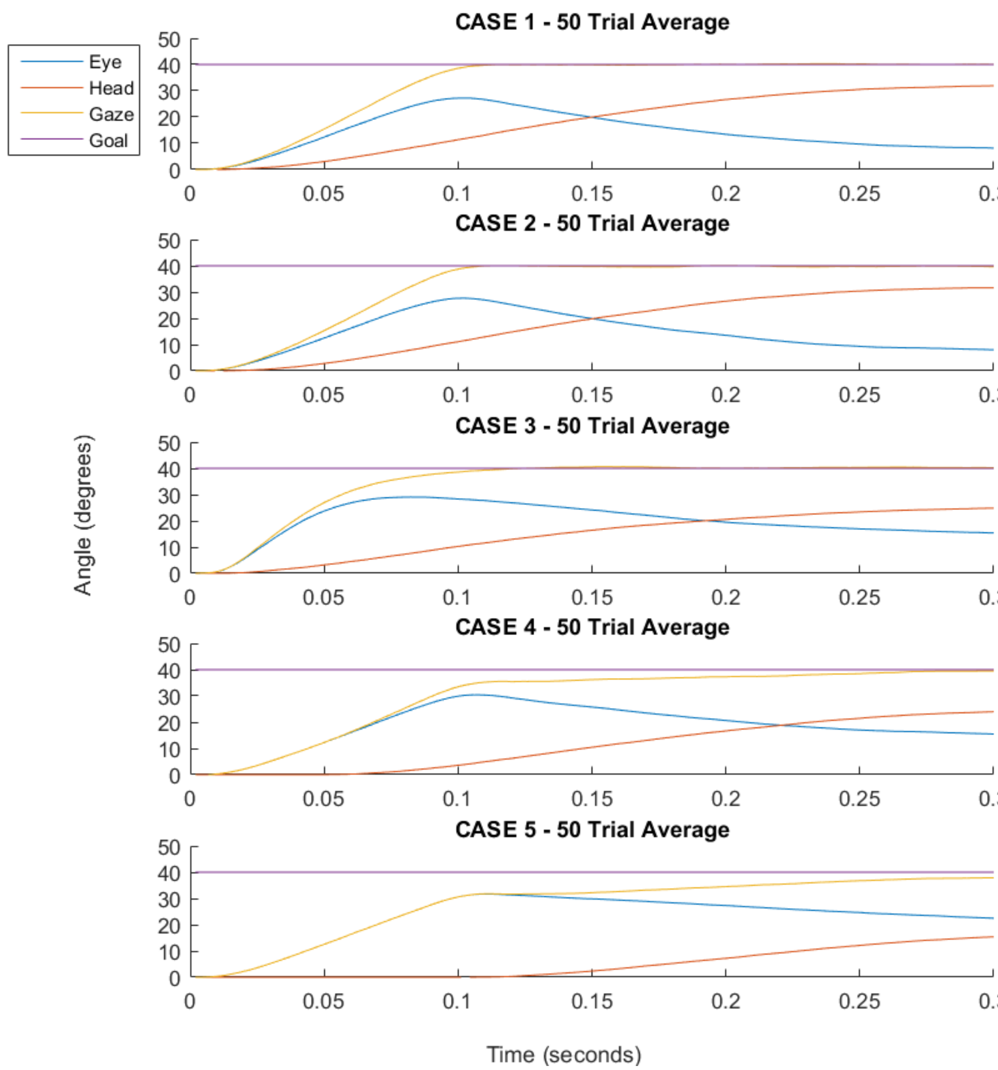

Discussion

The results of this simulation depict a base case (CASE 1, without signal-dependent noise or head holding) in which the gaze reaches the target position promptly at 100ms and maintains its position for remainder of the simulation's duration. Consistent with the natural gaze behavior, the eyes approach the target more rapidly than the head, reach a peak eccentricity, then more steadily return to their resting state while the head continues to turn in order to maintain the gaze.

Meanwhile, results of the cases with small and large signal-dependent noise (CASE 2 and CASE 3, respectively) demonstrate the ability of the computed optimal gains to effectively account for the noises in a way that ensures the overall gaze still reaches the target at 100 ms and holds its position for the remainder of the duration. The caveat appears to be a necessary increase in the rate at which the eyes move to their peak eccentricity and a decrease in the rate at which they return to their resting state - meanwhile, the head decreases the rate at which it turns altogether. This behavior is understood in observing that the head, being the more massive component of the system, requires significantly larger input signals (muscle activations) than the eyes would require to move the same distance. Since signal-dependent noise penalizes larger input signals, larger noises encourage a preference for eye movements over head movements to reach the same gaze positions.

Finally, the cases in which the head is held for a varying durations of the simulation (CASE 4 and CASE 5) exhibit a general inability of the computed optimal gains to accommodate for the inability of the head to complete eyes' work in directing gaze towards the target. More similar to the base case than the cases with signal-dependent noise, these cases exhibit the eyes moving to their peak eccentricity and returning to their base state at a relatively high rate. Meanwhile, the head attempts to direct the gaze towards the target as soon as it is able, but at a rate that appears to greater than that of the base case. In neither case (short duration or long duration of head holding) does the gaze seem to actually reach the target position in the duration of the simulation. This behavior is consistent with the fact that there is a relatively small noise associated with state updates in the system model, which means that the pose estimates generated from the Kalman gains will take a significant number of erroneous head-held predictions to adapt to the head holding - seemingly more than the number by which we wish to reach the target.