Challenges in Single-Camera Remote Eye Tracking

Martin Böhme and Erhardt Barth

Institute for Neuro- and Bioinformatics
University of Lübeck, Ratzeburger Allee 160, 23538 Lübeck, Germany

Introduction

Many eye tracking systems either require the user to keep their head still or involve cameras or other equipment mounted on the user's head. While acceptable for research applications, these limitations make the systems unsatisfactory for most AAC (Augmentative and Alternative Communication) applications.

So-called "remote" eye tracking systems, which allow the user to move their head freely within certain limits, have been available for a while. Traditionally, these systems have used cameras with long focal lengths to obtain a sufficiently high-resolution image of the eye. Because of the narrow field of view, the user's head movements must be compensated for, either by panning and tilting the camera itself or by using a movable mirror. This means that the head movement speed is limited by the speed with which the mechanical system can track the eye. Furthermore, once tracking is lost, reacquiring the eye is difficult because the camera has only a narrow field of view.

With recent increases in the resolution of CCD and CMOS cameras, it has become feasible to use fixed cameras with wide field of view for eye tracking. In this approach, the camera covers the whole area within which user's head may move while still imaging the eye with sufficient resolution for eye tracking.

One important task in remote eye tracking is measuring the position of the user's eyes. The most straightforward way of doing this is to use two or more cameras so that features in the images can be triangulated to determine their position in space. However, using two cameras instead of one increases the cost and complexity of the system substantially. Cost is of particular concern for AAC applications; in the following, we will therefore investigate the single-camera remote eye tracking problem.

While a number of researchers have proposed algorithms for calibrating single-camera remote eye trackers (1, 2), the results appear to be not as accurate as those achieved using fixed or head-mounted devices (0.5 to 1 degree of accuracy). Commercial remote eye trackers with high accuracy are available (3), but no implementation details have been published.

In this talk, we will report on our on-going work on calibration algorithms for remote eye trackers that aim to achieve an accuracy similar to that of fixed or head-mounted systems. Our results on simulated test data from an artificial eye model are quite promising, and we hope to achieve similar accuracy when we implement the algorithm on hardware in the near future.

Method

Most videographic eye trackers work by illuminating the eye with an infrared (IR) light source. This light source produces a glint on the cornea (the "corneal reflection" or "CR"), and the gaze angle is computed from the offset between the CR and the centre of the pupil using bilinear or biquadratic interpolation. The coefficients of the interpolation function are computed from data obtained during a calibration phase, during which the user is asked to fixate a certain number of points with known locations.

Our approach to remote eye tracking also uses infrared illumination, but instead of one light source, we use two. The distance between the CRs produced by these light sources can then be used to determine the distance of the eye from the eye tracker. This, together with the location of the eye in the camera image, allows us to deduce the three-dimensional position of the eye relative to the camera.

Note that we only determine the position and orientation of the eye; the position and orientation of the head are irrelevant for us since our approach does not use any reference points on the head.

Using an interpolation scheme to calculate gaze position from the observed pupil and CR positions, as for the fixed-head eye tracker, does not appear to be an option for the remote eye tracking scenario because it has a far greater number of degrees of freedom – covering the whole space of possible eye positions and eye orientations during calibration would not be feasible.

We therefore believe that a calibration procedure for remote eye tracking must be based on a model of all relevant physical properties of the human eye. Of course, the shape and size of the eye vary from person to person, so the model must contain a suitable set of parameters to accommodate these differences. Calibration then means estimating the values of these parameters for a specific person.

To date, our eye model contains the following parameters:

- r_{cornea}: The radius of curvature of the corneal surface (which we assume to be spherical)

- r_{pc}: The distance between the centre of corneal curvature and the pupil centre

- α_{fovea}: The angular offset between the optical axis of the eye and the direction of gaze, which caused by the fact that the fovea does not lie on the optical axis but is offset temporally and slightly upwards (at the moment, we only model the horizontal component of this offset).

The values of these parameters for a particular user are determined by taking the pupil and CR positions for a set of calibration points and then varying the parameter values to minimize the error between the observations predicted by the model and the actual observations.

Results and Outlook

We implemented our calibration algorithm in Matlab and assessed its performance on simulated test data. For the tests, the user was assumed to be seated at a distance of 50 cm from a 40x30 cm screen. To evaluate the robustness of our approach to noise, we added a certain amount of random error to the measurements of pupil centre and CR position.

The results we have obtained so far are encouraging: Assuming a maximum measurement error of 0.5 pixels, the maximum error in gaze position is 13 mm (1.5 degrees), with an average error of 4.5 mm (0.5 degrees). An error of this magnitude should be more than acceptable for most AAC applications. The assumed measurement error should be achievable if some care is taken in the image processing and camera calibration steps.

The next step, then, is to implement our algorithm on actual hardware. This will reveal whether the algorithm can live up to the potential it has demonstrated on our simulated data.

References

(1) C. H. Morimoto, A. Amir and M. Flickner. Detecting Eye Position and Gaze from a Single Camera and 2 Light Sources. In: *16th International Conference on Pattern Recognition*, 314–317 (2002).

(2) T. Ohno, N. Mukawa. A Free-head, Simple Calibration, Gaze Tracking System That Enables Gaze-Based Interaction. In: *Eye Tracking Research and Applications 2004*, 115–122.

(3) Tobii 1750 eye tracker, Tobii Technology AB, Stockholm, Sweden.