INFERRING EYE MOVEMENTS ON THE BASIS OF HEAD AND VISUAL TARGET POSITION

Lynn M. Batten[1] and Bill Whiten[2]

The modern optical lens design process relies on ray-tracing that requires the geometry of the visual task to be simulated in order to derive the optical parameters. Such measurements must take into account distant, intermediate and near visual tasks. SOLA has developed a non-intrusive, low cost system which tracks eye movements during the reading process. This report analyses their procedure in order to determine the accuracy of the tracking system. It concludes that some assumptions in the existing algorithm are overly restrictive, while, overall, this is an effective tracking method. In addition, a Fourier analysis of the sampling rate demonstrates that 10 Hz is a sufficiently high rate to use, and that lossy compression is adequate for their needs.

1. Description of the procedure

The optical lens design process relies increasingly on ray-tracing that requires the geometry of the visual task to be simulated in order for the optical parameters to be calculated and plotted. In a research setting, there are several options for making such measurements. Some are intrusive; some are not. Most are expensive. (See [2], [4], [5]. A recent conference on the subject of eye-tracking is [6].) Recently, SOLA developed a non-intrusive, low cost eye-tracking system based on using sensors to track head and book positions.

The bulk of this section, along with the figures, is from Sola's summary of the project as provided to MISG.

Apparatus and procedure

Subjects wear a near or intermediate prescription, depending on the task, fitted to a custom lens clip attached to a demonstrator frame. (See Figures 1 and 2.) The clip is fitted with a pair of stock single vision lenses that include

[1]School of Computing and Mathematics, Deakin University, 662 Blackburn road, Clayton VIC 3168, Australia. Email lmbatten@Deakin.edu.au

[2]Julius Kruttschnitt Mineral Research Centre, The University of Queensland, Isles Road, Indooroopilly, Brisbane QLD 4068, Australia. Email W.Whiten@acm.org

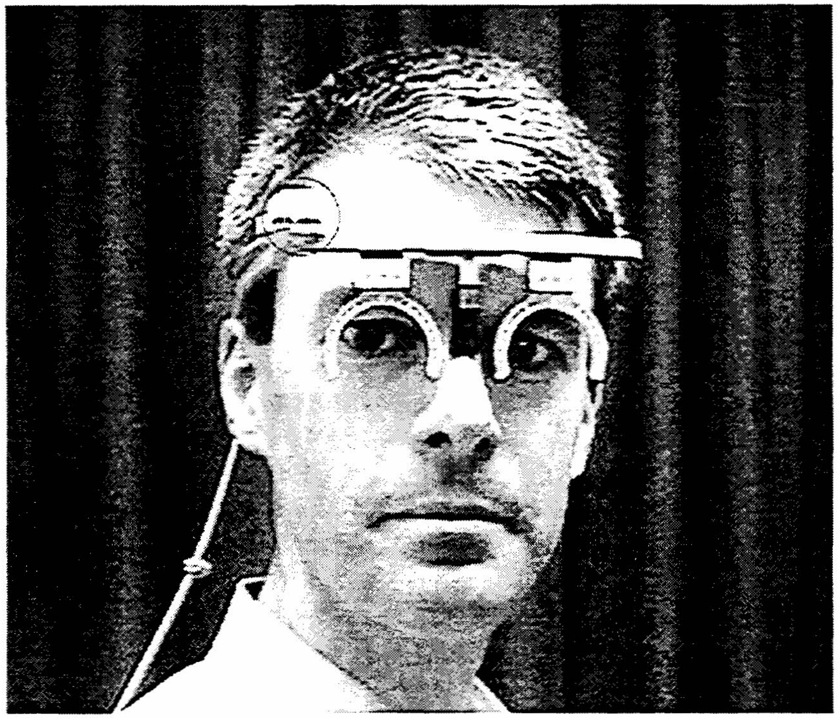

Figure 1: Subject wearing the standardised reading lenses and frame. A Progressive Addition Lens (PAL) demonstrator frame was used to mount custom made Single Vision near prescription and intermediate prescription lenses. The model in this photograph is not wearing the clip in lenses. A Polhemus sensor cube (in circle) was mounted to the right frame temple to monitor head movements.

the near sphere and Cyl correction. The lens is circular, 38 mm diameter. Also attached to the frame is a Polhemus electromagnetic motion sensing cube (see Figure 1). The average distance position of the head receiver cube and the apex of the cornea of the right eye is 30 mm temporal, 30 mm above and 20 mm in front of the corneal apex.

While wearers perform a standardised near and intermediate reading task, the position of both the head and book are recorded (10 Hz sampling rate) in X, Y, Z, Azimuth (head turn, t), Elevation (head nod angle, v) and Roll space (side to side tipping, r) using a Polhemus electromagnetic motion sensor system.

Near and intermediate tasks

Subjects read a standardized near and intermediate stimulus as described below. During these tasks, head and stimulus position are recorded in real-time with a 10 Hz sampling rate and to an accuracy of approximately 5 mm (X, Y and

Figure 2: Transmitter and receiver cubes. In the configuration shown there is one transmitter cube shown on the stand to the right in the image (large circle) and two receiver cubes, one on the wearer's head attached to the trial-frame and one inside the clip board (small circles).

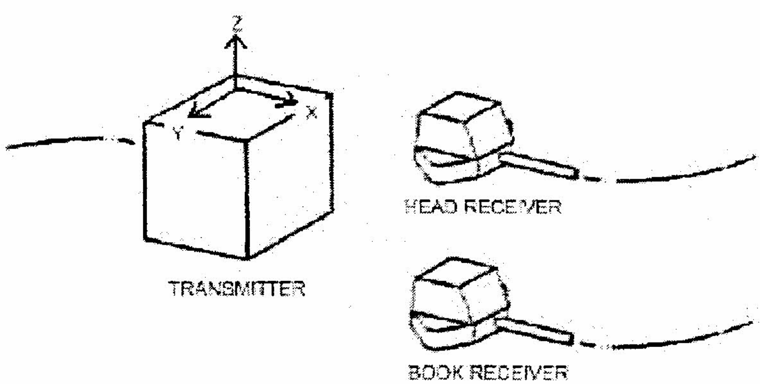

Figure 3: Orientation and position of the transmitter and receiver cubes. The z axis is negative in the direction shown while all other axes are positive. The receivers have the same coordinate system.

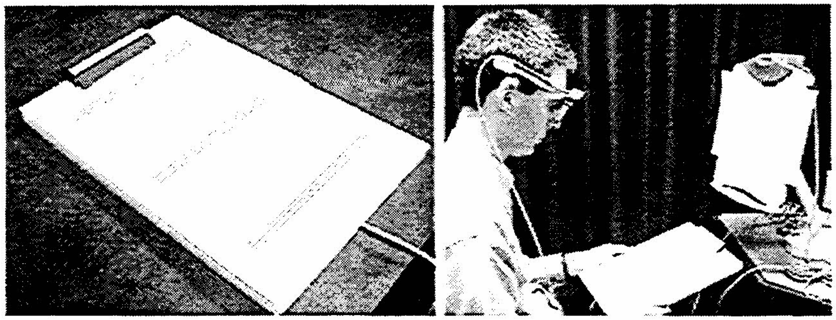

Figure 4: Standardised near reading stimulus and layout. Clip-board (left) and subject reading (right).

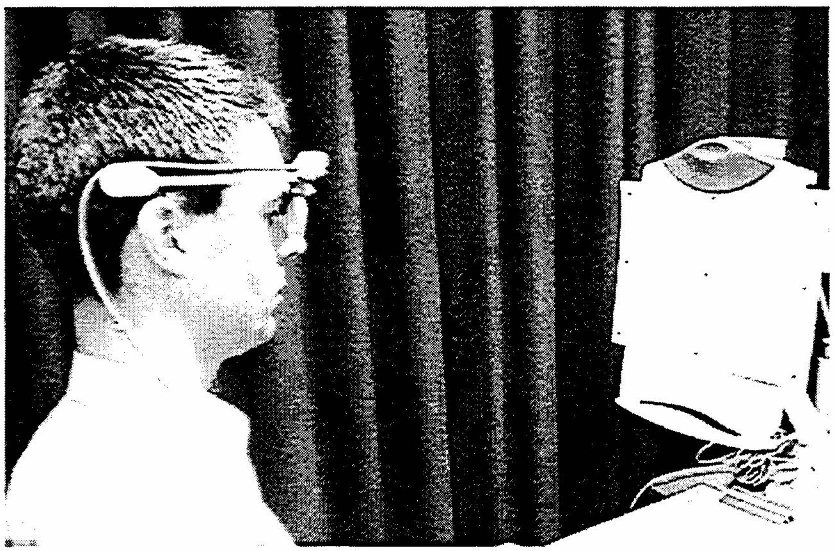

Figure 5: Intermediate stimulus and layout. Three rows of 10 point font Times New Roman three digit numbers are printed in 'Landscape' mode.

Z) and 1 degree (Azimuth elevation and roll). The near reading task requires the subject to read text broken into three paragraphs in 10 point Times New Roman font (see Figure 4).

Three paragraphs of three line 10 point font, Times New Roman letters are printed so that the top and bottom paragraphs are 120 mm above and below the middle paragraph. The paragraphs are 170 mm wide. The text is placed on a clip board with a sensor cube imbedded in the board to be coincident with the centre of the middle paragraph. Subjects read silently. The subject in Figure 4 is pictured reading the book on their lap, but allowance is made for reading on the desk also. The intermediate reading target is printed on an A4 page (210 mm by 297 mm) in 'landscape' mode. Three rows of three digit numbers in 10 point 'Times New Roman' font are printed so that the top row is 15 mm from the top of the page, the second row is 90 mm below the first and the final row 90 mm below that. The columns flanking the centre column are 135 mm to the left and right of centre. The target is clipped onto the transmitter stand so that the middle column of numbers is in line with the centre of the transmitter cube. For the intermediate data, Sola are mainly interested in eye-turn. Eye elevation is not required.

For the purpose of clarity in this illustration the 3 digit numbers are superimposed with black spots. Subjects read the numbers aloud from the top left to bottom right in rows. The intermediate stimulus/transmitter stand is arranged so that the centre number on the top row is at eye-level and approximately 65 cm from the wearer's eyes.

In summary, the system described above allows us to measure:

- head and book position in three dimensions relative to a fixed sensor;
- the turn, elevation and roll of the head and book.

From these measurements the reading distance from the head to the book can be calculated.

In addition, the following measurements are estimated:

- eye turn and elevation.

The questions tackled by the group include:

1. Are the eye turn and elevation estimates best possible?
2. Can readers be classified into a few distinct categories on the basis of their measurements?
3. Is the currently used sampling rate appropriate?

Our conclusions are listed in the last section of the report.

2. The geometry

Coordinates

The equipment determines the position and rotation of the two sensors relative to the fixed transmitter. The position coordinates (x, y, z) are permanently calibrated, however the rotations of the sensor can be zeroed with the sensor in any attitude to correspond to the axes defined by the transmitter. In the Sola application this is done before each test. To be able to relate the local head and book coordinates to the global coordinates, it is necessary that the relation between the physical local coordinates and those being set by the sensor calibration are known. In the case of the book this is easiest done by aligning the physical coordinates of the book (length, width and depth) with those of the transmitter. A small guide on a horizontal shelf could conveniently be used for this purpose. In the case of the head it is more difficult, as there is no naturally defined coordinate system for the head. This limits the accuracy to which eye rotations relative to the head can be calculated. Currently, for head axes calibration, the subject is told to look straight ahead with head vertical and straight. This should be sufficient for relative eye movements, but is limited for eye relative to head measurements. More accurate measurement of the eye to head angle would require a separate measure of eye position during the calibration. This could be done with a video camera facing the subject and image processing.

The equipment uses Euler coordinates [7] to describe the rotation of the sensors relative to the axes of the fixed transmittor. These coordinates are measures of the rotations needed to move from the fixed (x, y, z) coordinates to those of the sensor [7]. Three angles are required and these are obtained as rotations about each axis in the local coordinates as these are rotated to the observed position. Starting with the local coordinates aligned with the fixed coordinates, the first rotation is about the z axis (turn or azimuth t), the second rotation is about the new y axis (elevation v which is negative for a rotation upward to preserve the convention for direction of rotations), and the final rotation is about the latest position of the x axis (roll r). These three angles uniquely define the directions of the sensor based coordinates.

A sample of a typical recording of the head position (x, y and z) and angles (t, v and r) is given in Figure 6, and the corresponding recordings for the book in Figure 7.

Calculation of position

The position of the text relative to the book sensor, and the eye relative to the head sensor are, since they are fixed relative to the sensors, measured in the

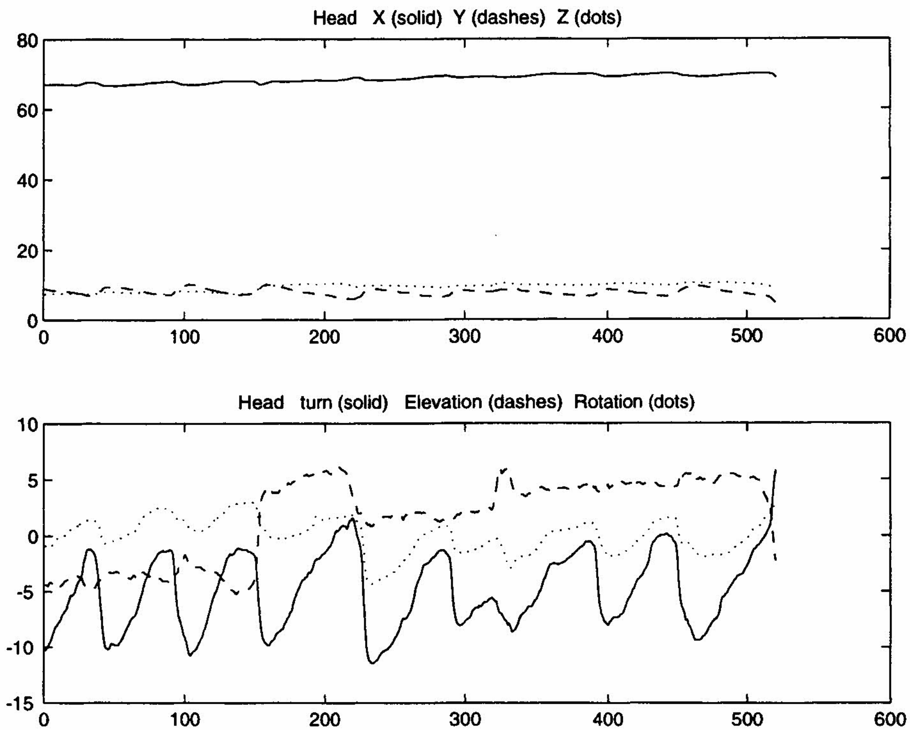

Figure 6: Sample head position and angle data. X axis is sample number (0.1 seconds), and Y axis top is com, bottom is degrees.

local coordinates of the sensor. Neither of these is at the actual position of the sensor. To determine the relative position of the eye and reading position in the book, it is useful to convert to a single coordinate system. This can be done by first converting to absolute coordinates, namely those of the transmitter. It is then easy to calculate the relative position in the absolute coordinate frame. This position can then be converted to the local coordinate frame of the head; then the distance to the text being read and the angles of the eye (turn and elevation) relative to the head can be calculated. This procedure can be summarized as:

1. Convert the local eye coordinates relative to the head sensor to coordinates parallel to the fixed axis of the transmitter and then to the actual transmitter axes.

2. Convert the local text coordinates relative to the book transmitter to coordinates parallel to the fixed axes of the transmitter and then to the actual transmitter axes.

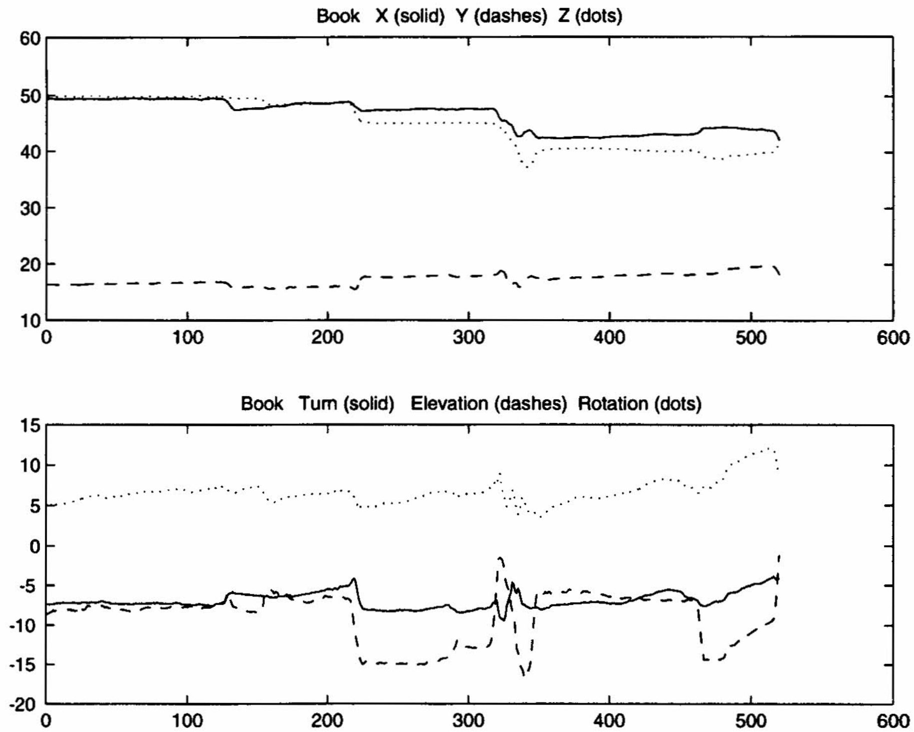

Figure 7: Sample book position and angle data. X axis is sample number (0.1 seconds), and Y axis top is com, bottom is degrees.

3. Convert the absolute positions of the eye and text to positions relative to the eye in coordinates parallel to the fixed coordinates.

4. Convert the position of the book relative to the eye to the axes of the current head position.

5. Convert the local (x, y, z) coordinates of the text relative to the eye in the head coordinates to distance, turn, and elevation.

6. Finally as the local head x axis is into the face, convert the calculated turn angle (which is close to 180 degrees) to one relative to the angle from the axes with the x axis facing forward.

The details of the required rotations are as follows.

To convert a point (x^*, y^*, z^*) given in the sensor coordinates, centred on the eye and parallel to the head, to an absolute position using the three Euler angles of the sensor coordinate frame given by the three Euler angles $(t,\ v,\ r)$, requires that the point position be calculated in the new axes that are created by rotating through the negative of the Euler angles in the reverse order to which the angles were defined to reach the sensor coordinate frame [1].

More explicitly, we can describe each matrix operation separately as follows where x, y and z are the coordinates relative to the sensor in axes parallel to the fixed coordinates of the transmittor. As the origin of these coordinates is known with respect to the transmittor coordinate frame, a simple addition then gives the position of the point (x^*, y^*, z^*) in the fixed coordinate frame.

Undo roll about x^* axis (rotation is positive y to z):

$$\begin{bmatrix} x_1 \\ y_1 \\ z_1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(r) & -\sin(r) \\ 0 & \sin(r) & \cos(r) \end{bmatrix} \begin{bmatrix} x^* \\ y^* \\ z^* \end{bmatrix}.$$

Undo elevation about y_1 axis (rotation is positive z to x):

$$\begin{bmatrix} x_2 \\ y_2 \\ z_2 \end{bmatrix} = \begin{bmatrix} \cos(v) & 0 & \sin(v) \\ 0 & 1 & 0 \\ -\sin(v) & 0 & \cos(v) \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \\ z_1 \end{bmatrix}.$$

Undo turn (azimuth) about z_2 axis (rotation is positive x to y):

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} \cos(t) & -\sin(t) & 0 \\ \sin(t) & \cos(t) & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_2 \\ y_2 \\ z_2 \end{bmatrix}.$$

This can be compressed to:

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} \cos t & -\sin t & 0 \\ \sin t & \cos t & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos v & 0 & \sin v \\ 0 & 1 & 0 \\ -\sin v & 0 & \cos v \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos r & -\sin r \\ 0 & \sin r & \cos r \end{bmatrix} \begin{bmatrix} x^* \\ y^* \\ z^* \end{bmatrix}.$$

Given the two positions in the fixed frame of the eye and text, the relative position of the text is calculated as the difference (text − eye). This is then a position relative to the eye which needs to be converted to coordinates parallel to those of the head sensor (note these are not the same coordinates as those of the head sensor as the origin is different). If the relative coordinates in the axes parallel to the transmittor frame are (x, y, z) then the rotations required are the inverse of the rotations used above. Note that the inverse rotations must be applied in the reverse order to the previous calculation, namely:

$$\begin{bmatrix} x_2 \\ y_2 \\ z_2 \end{bmatrix} = \begin{bmatrix} \cos(t) & \sin(t) & 0 \\ -\sin(t) & \cos(t) & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix}$$

$$\begin{bmatrix} x_1 \\ y_1 \\ z_1 \end{bmatrix} = \begin{bmatrix} \cos(v) & 0 & -\sin(v) \\ 0 & 1 & 0 \\ \sin(v) & 0 & \cos(v) \end{bmatrix} \begin{bmatrix} x_2 \\ y_2 \\ z_2 \end{bmatrix}$$

$$\begin{bmatrix} x^* \\ y^* \\ z^* \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(r) & \sin(r) \\ 0 & -\sin(r) & \cos(r) \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \\ x_1 \end{bmatrix}.$$

Finally the head spherical polar coordinates (distance, turn, elevation) are calculated as:

$$\begin{aligned} d &= \sqrt{x^{*2} + y^{*2} + z^{*2}} = \sqrt{x^2 + y^2 + z^2} \\ t &= \arctan(y^*/x^*) \\ v &= \arctan(z^*/\sqrt{x^{*^2} + y^{*^2}}). \end{aligned}$$

This gives the required eye motion information for reading the text at the position given.

3. Coding the geometry

Existing Sola code is based on the instructions in the manual for the tracking equipment used. The MISG project found some discrepancies between the Sola code and the manual. In particular, the Sola code assumes that roll is measured about the fixed x-axis with respect to the transmitter, whereas according to the manual, the machine measures this with respect to the sensor's x-axis [Manual b–12]:

> *Euler angles are defined as the sequence of angles (azimuth, elevation, and roll) that define the orientation of the sensor with respect to the X, Y, Z alignment reference frame. Azimuth is a rotation of the sensor's x-axis projection in the X Y reference plane about the Z reference axis. Elevation is a rotation of the sensor's x axis about the Y reference axis. Roll is a rotation of the sensor's y (or z) axis about its x axis.*

Other, not necessarily appropriate, restrictions set into the Sola code include an assumption that the book and the person reading it are aligned in the same plane, and that the vector from the eye to the centre of the book is at right angles to the book. New code was developed by the reporting team based on the following algorithm.

MISG Algorithm for Processing of Spatial Data

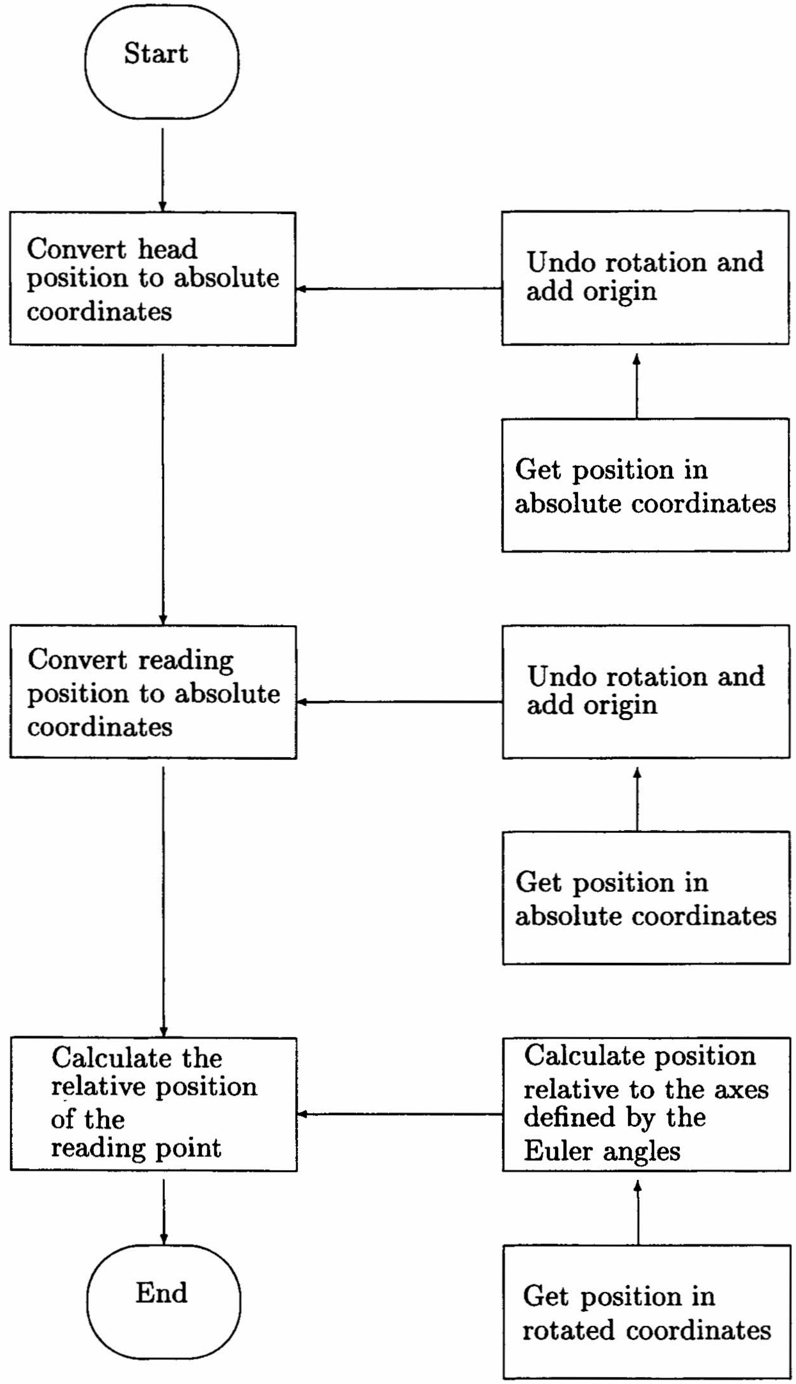

Comparison of Sola's code with our alternative code

		Our code			Sola's code
Paragraph	mean/ stand. dev.	D	T	E	E
1^{st}	m	47.83	-8.17	42.45	42.54
	sd	0.17	1.21	0.96	
2^{nd}	m	47.43	-8.60	44.83	45.32
	sd	0.26	1.36	0.19	
3^{rd}	m	46.88	-9.38	48.77	49.46
	sd	0.30	1.75	0.58	

Table 1: Data when looking straight ahead.

		Our code			Sola's code
Paragraph	mean/ stand. dev.	D	T	E	E
1^{st}	m	51.13	-4.93	27.10	29.21
	sd	0.29	0.82	0.79	
2^{nd}	m	49.30	-7.41	32.03	39.02
	sd	0.75	2.59	1.28	
3^{rd}	m	52.62	-11.77	38.73	47.21
	sd	0.63	4.10	1.37	

Table 2: Data when looking to the side.

We established that the SOLA routine projected all the data onto the x-z plane (looking straight ahead). This turned out to be an overly optimistic assumption.

4. Separation of sections

The actual recordings of the near data include reading from three positions on the page and the transitions to and from these sections. While it is in principle possible to read without head or book motion, normal reading includes both head and book motion as well as eye motion. Using these motions, it is normally possible to recognise the section of the record where the subject is actually reading the text. However there are considerable differences in the manner in which people read; for instance some readers may move the head while holding the book still, some may move the book and keep the head still, while others may move both, and yet others, neither.

The range of different reading behaviours makes the construction of a simple algorithm for the recognition of the three sections difficult. What is readily seen in the data is that there are much larger motions in some of the variables during the sections between reading the text than when reading it. This indicates that a derivative of the measurement would be useful in the discrimination of the reading sections. The lower half of Figure 8 shows the changes between the successive 0.1 second values of head relative position of the book, that are given in the upper half of Figure 8. As there is a large variation between different people, several different variables may need to be examined and the ones that give the best discrimination in the particular case used. Experience in processing the data is needed to determine which variables are most likely to be useful.

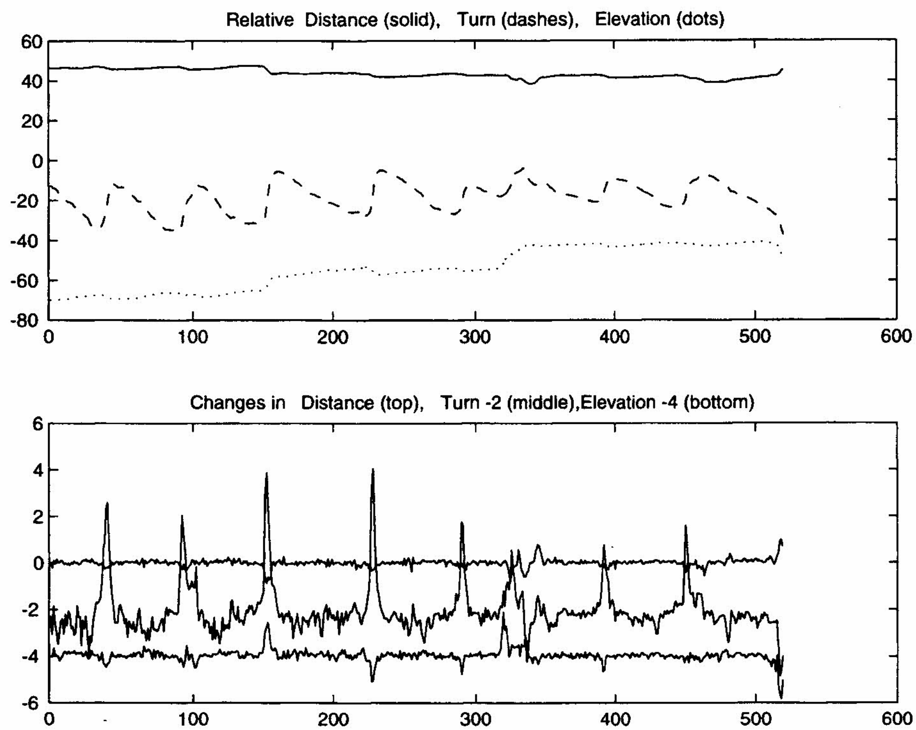

Figure 8: Relative position data (head to book) and change between successive sample value (bottom). X axis is sample number (0.1 seconds), and Y axis top is cm and degrees, bottom is change in cm and degrees.

Once a reasonable estimate of the reading sections has been developed, an evaluation of the most useful variable can be done by comparing the ratio of the standard deviations in the reading section with those in the transition sections (probably with the initial and final sections removed). A ratio significantly different from the value one indicates that the variable is useful for discrimination.

However it needs to be kept in mind that the same variable may not be useful in all cases, and in fact probably two or three different variables may need to be used.

A manual examination of twenty data sets from the near reading task gave a separation of the sequences into portions corresponding to the reading of the three paragraphs in 18 cases. These were used to examine the amount of separation of the reading section from the between reading sections. The root mean square of the ratio of measures of variation between the reading and between sections has been calculated. Three measures were used:

- the standard deviation of the sections;
- the standard deviation of the change between successive data points; and
- the standard deviation of the residues after a linear trend has been removed.

The preceding three are combined as a root mean square to give an overall value.

These were calculated for each of 15 measurements (6 head position, 6 book position, 3 relative position) to give Table 3.

Variable		Value	Deriv.	Detrend.	RMS
Head	x	0.56	0.64	0.58	0.59
	y	0.89	0.81	0.80	0.83
	z	0.34	0.61	0.34	0.45
	t	0.92	0.82	0.85	0.86
	v	0.28	0.48	0.30	0.36
	r	0.87	0.78	0.88	0.84
Book	x	0.42	0.53	0.49	0.48
	y	0.72	0.66	0.58	0.66
	z	0.85	0.53	0.74	0.72
	t	0.54	0.56	0.57	0.56
	v	0.55	0.51	0.45	0.51
	r	0.65	0.53	0.62	0.60
Rel.	d	0.50	0.55	0.48	0.51
	t	0.79	0.84	0.77	0.80
	v	0.23	0.41	0.23	0.30

Table 3: Relative variation reading to non-reading (i.e. changing to next paragraph).

It can be seen from Table 3 that relative elevation (v) is giving the best relative change between reading and not reading. Other values that give low ratios are head x z and t, book x and t, and relative d. Book z does not give a good ratio overall, but in two cases gives a low value when relative elevation (v) is high.

As might have been expected y, turn, and rotation, do not give good ratios for distinguishing between reading and not reading. The head turn sometimes shows the line reading of the 3 lines in each paragraph, as can be seen in Figure 8.

An estimate of which sections correspond to reading can be attempted using a trimmed mean of the difference between sucessive readings [3]. This involves an iterative calculation of the mean and standard deviation at each stage removing points that are more than say two standard deviations from the mean. This will, at least for most of the measurements, allow us to separate reading sections from the transition sections. Convergence can be assisted by good initial estimates of either the mean and standard deviation, or of the points that correspond to reading. An initial smoothing of the data could help.

The question arises as to whether a combination of variables would be more useful than a single variable. For a linear combination of a small number of variables we could try searching for the combination that gives the largest ratio of variances.

To automatically determine reading sections a suitable procedure could be based on first identifying possible reading sections by dividing the data according to the expected location of the reading sections (1 below), finding regions of low standard deviation (2 to 7), then iteratively fitting a line to the region identified (8), and extending the line to cover points that follow the trend of the line (9, 10 and 11). The steps needed to do this are:

1. Divide the recording into three equal sections which should roughly correspond to reading the three separate paragraphs, and repeat the following for each section.

2. Within the section determine subsections that have a low standard deviation by dividing each section into say ten near equal subsections and calculating the standard deviation for each interior subsection.

3. Choose the three subsections that have the lowest standard deviation as the base interval for further calculations.

4. Calculate the mean and standard deviation for this base set.

5. Find all the points in the interval that differ from this mean by less than say three standard deviations.

6. Apply a smoothing over five points to eliminate isolated values.

7. Choose the longest region of points with low standard deviation as the start for the selection of a detrended interval.

8. Calculate the trend line of the selected interval, and the residuals for the whole data.

9. Apply a three point smoothing to the residuals.

10. Calculate the standard deviation of the smoothed residuals over the selected interval.

11. Increase or decrease the interval progressively at each end, to include points that have a residual of less than say three standard deviations and omit those with greater than three standard deviations

12. Repeat steps 8 to 11 until there are no further changes.

Testing this on twenty samples using the relative elevation gave 18 reasonable results, and one case in which reading had apparently started late and was acceptable if some of the initial data was omitted. The final case has defied manual separation of the reading sections.

Both the failure cases are easily identified by overlapping of the calculated reading intervals which could be used as an indication that further analysis is required.

An example of the result of this procedure is given in Figure 9.

As there is wide variation of the reading behaviour of different people it is suggested that a graph showing the recognized reading periods be available for the operator to check. A recording of the person's voice during the test would be useful in resolving difficult cases.

5. Processing the data – the Fourier analysis

A Fourier analysis was conducted on several of the recordings. It was found that only the low fequency terms contributed to the recording and a large number of terms could be dropped without significantly degrading the signal. The result is a smoothed record that requires significantly less information to store and as

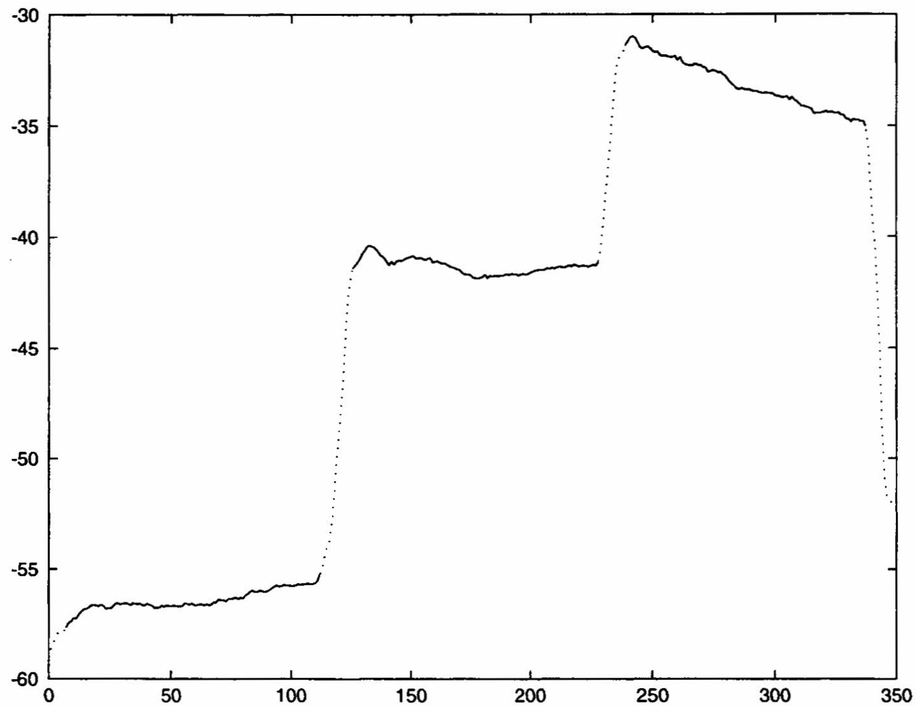

Figure 9: Identification of reading sections using relative elevation angle. Solid line indicates reading. X axis is sample number (0.1 seconds), and Y axis degrees.

a considerable amount of noise has been removed could also be easier to process to determine the three regions where the subject is actually reading.

Figure 10 shows the effect of retaining only 32%, 24%, and 16% lower frequency Fourier coefficients. It can be seen that little information is lost after removing 68% of the coefficients. This roughly corresponds to dividing the data into groups of three points and retaining the average values of these groups.

The fact that only low frequency components appear in the data indicates that there is no need to change to recording at a higher frequency. Some effects seem to appear in the data at about 0.5 second intervals, hence it is concluded that the current recording at 0.1 second intervals gives an adequate but not excessive margin over the variation in the motion being recorded. Recording at a lower frequency may have a detrimental effect on the current method of dividing the signal into three reading sections.

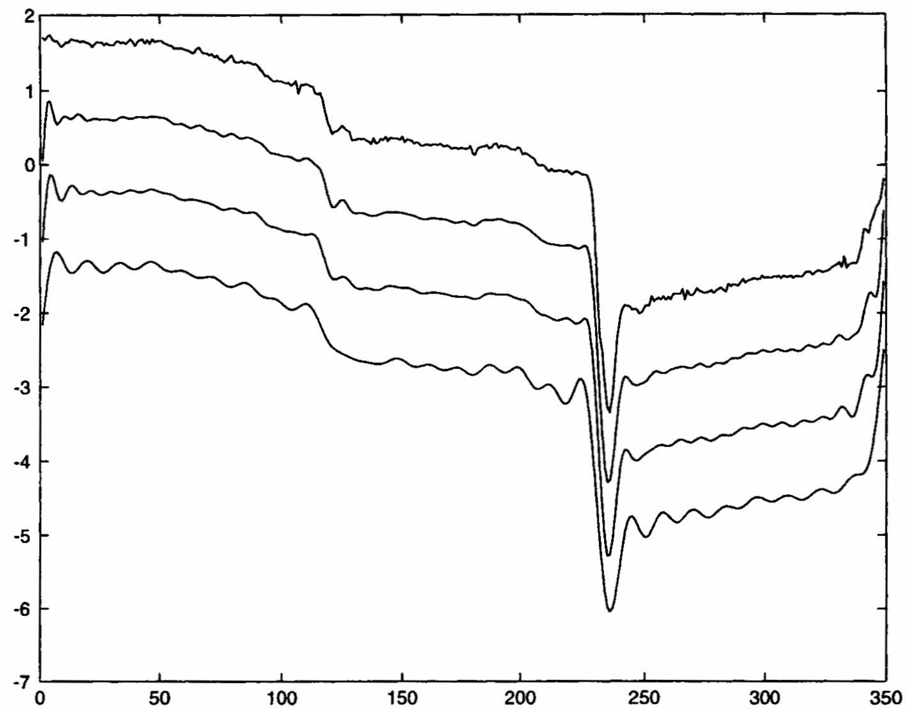

Figure 10: Original signal (book x coordinate, top) and truncated Fourier reconstructions (progressively offset) using 32%, 24%, and 16% of the Fourier coefficients. X axis is sample number (0.1 seconds), and Y axis is cm with shifted origins.

6. Conclusions.

This project spent a large proportion of the time on the basic processing of the data. The program used by Sola was evaluated in detail and, independently, a program doing the basic calculations from first principles was developed. A comparison of the two indicated that the Sola program lost accuracy when the subject was not aligned with the x axis of the transmitter.

It was found that the separation of the sections where the subject was actually reading could be based on the relative elevation of the book.

An algorithm based on identifying linear segments with a relative low standard deviation s appears suitable for identifying the reading sections.

A Fourier analysis showed that the recording consists of only lower frequency components. This then led to the following conclusions:

- The data could be compressed without significant loss of accuracy by discarding the higher frequency components.
- A Fourier analysis may be useful for removing noise before other processing is undertaken.
- The frequency of data recording is quite sufficient to obtain all the useful signal that is available.

Acknowledgements

The project moderators, Lynn Batten and Bill Whiten, are grateful to Scott Fisher and John Bonnett from Sola for presenting the problem to the MISG and support during the week. Cooperating with the industry representatives on this project was a team of enthusiastic individuals with various complementary backgrounds: Yelena Davydova, Kathrine Grillett, Darren Lohmeyer, Stephen Lucas, Scott Mackenzie, Julia Piantadosi, Jane Thredgold, Bruce van-Brunt.

References

[1] R.O. Hill, *Elementary Linear Algebra with Applications*, Third Edition, (Saunders College Publishing, New York, Sydney, Tokyo, 1996).

[2] C. Lee, "Eye and head coordination in reading: roles of head movement and cognitive control", *Vision Research* **39** (1999), 3761–3768.

[3] S. Rasmussen, *Introduction to Statistics with Data analysis* (Brooks/Cole Publishing, Pacific Grove, California, 1992).

[4] K. Rayner, "Eye movements and cognitive processes in reading, visual search, and scene perception", *Eye Movement Research*, Findlay *et al.* Editors, 1995, 3–17.

[5] E.D. Reichle, A. Pollatsek, D. Fisher and K. Rayner, "Toward a model of eye movement control in reading", *Psychological Review* **105** (1998), 125–157.

[6] S.N. Spencer, *Proceedings Eyes Tracking Research & Applications Symposium 2000* (ACM, New York, 2000).

[7] C.E. Weatherburn, *Elementary Vector Analysis* (G. Bell and Sons Ltd., London, 1967).