



Research White Paper

WHP 165

July 2008

A multi-cue bayesian state estimator for gaze prediction in open signed video

S.J.C. Davies, D. Agrafiotis, C.N. Canagarajah, and D.R. Bull
Department of Electrical and Electronic Engineering, University of Bristol, Bristol, BS8 1UB, UK

BRITISH BROADCASTING CORPORATION

**A multi-cue bayesian state estimator
for gaze prediction in open signed video**

S.J.C. Davies, D. Agrafiotis, C.N. Canagarajah, *and* D.R. Bull
Department of Electrical and Electronic Engineering, University of Bristol, Bristol, BS8 1UB, UK

Abstract

On most delivery platforms, bandwidth or bit-rate is limited; it thus becomes important to match the quality of compression coding used for any service to the particular service needs. An approach to compression which preferentially assigns values of precedence to the “salient” elements of a scene offers a potential route to more efficient or successful coding especially in bit-rate scarce delivery systems. This paper describes research on optimising the coding of open-signed video which has been undertaken by Sam Davies at the University of Bristol as part of his BBC-sponsored Ph.D. The work has direct potential relevance to the provision of signed TV content but will have other potential applications relevant to a broadcaster.

Additional key words: none

White Papers are distributed freely on request.
Authorisation of the Head of Broadcast/FM Research is
required for publication.

© BBC 2008. All rights reserved. Except as provided below, no part of this document may be reproduced in any material form (including photocopying or storing it in any medium by electronic means) without the prior written permission of BBC Future Media & Technology except in accordance with the provisions of the (UK) Copyright, Designs and Patents Act 1988.

The BBC grants permission to individuals and organisations to make copies of the entire document (including this copyright notice) for their own internal use. No copies of this document may be published, distributed or made available to third parties whether by paper, electronic or other means without the BBC's prior written permission. Where necessary, third parties should be directed to the relevant page on BBC's website at <http://www.bbc.co.uk/rd/pubs/whp> for a copy of this document.

**A multi-cue bayesian state estimator
for gaze prediction in open signed video**

S.J.C. Davies, D. Agrafiotis, C.N. Canagarajah, and D.R. Bull
Department of Electrical and Electronic Engineering, University of Bristol, Bristol, BS8 1UB, UK

Foreword

UK broadcasters have a duty to make a proportion of their content accessible to those who are profoundly deaf by providing programmes which are sign-language interpreted or sign-language presented. Currently the BBC's provides over 5% of its output in signed form on the broadcast versions¹ of BBC ONE, BBC TWO, BBC THREE, BBC FOUR, CBBC, CBeebies & BBC NEWS.

On most delivery platforms, bandwidth or bit-rate is limited; it thus becomes important to match the quality of compression coding used for any service to the particular service needs. Vision coding typically already uses a number of generic strategies (including adaptive quantisation, motion-compensation) automatically to adapt the coding to the complexity of the image to be coded and to the available resource. An approach which also preferentially assigns values of precedence to the "salient" elements of a scene offers a potential route to more efficient or successful coding especially in bit-rate scarce delivery systems.

A crucial determinant of service quality for signed content is the "legibility" and hence comprehensibility of the sign-language as delivered to the user. Prior work based on eye-gaze tracking (and referenced in this paper) has shown that for a deaf viewer who can understand the signing the most important part of the image is understandably associated with the in-vision signer. The challenge is to find an appropriate dynamic balance between the coded image quality of the basic scene content and that of the signer to maintain this essential "legibility".

This paper describes research on optimising the coding of open-signed video which has been undertaken by Sam Davies at the University of Bristol as part of his BBC-sponsored Ph.D. The work has direct potential relevance to the provision of signed TV content but will have other potential applications relevant to a broadcaster.

Nick Tanton

BBC R&I

¹ This is as "open signing" where all viewers of a signed programme see the signer as part of the picture. The advent of ipTV and of hybrid receivers (eg. with broadband plus one or more other delivery platform) offers the chance of delivering signed content as an elective video alternative to its unsigned counterpart.

A Multi-Cue Bayesian State Estimator for Gaze Prediction in Open Signed Video

S.J.C. Davies*, D. Agrafiotis, C.N. Canagarajah, and D.R. Bull

July 1, 2008

Abstract

We propose a multi-cue gaze prediction framework for open signed video content, the benefits of which include coding gains without loss of perceived quality. We investigate which cues are relevant for gaze prediction and find that shot changes, facial orientation of the signer and face locations are the most useful. We then design a face orientation tracker based upon grid-based likelihood ratio trackers, using profile and frontal face detections. These cues are combined using a grid-based Bayesian state estimation algorithm to form a probability surface for each frame. We find that this gaze predictor outperforms a static gaze prediction and one based on face locations within the frame.

1 Introduction

The ability to predict where viewers of a video are looking at any given time would be a useful tool in many video coding scenarios - from bit rate allocation to error protection. Studies have shown that eye gaze (i.e. the point on which a human's eyes are focussing) is both a top-down and bottom-up guided process [14]. Previous work on gaze prediction has generally been non-context specific [15] [4] and therefore a largely bottom up process. Many techniques [5] [6] [17] involve the bottom-up concept of saliency [12] [16], which combines low level features of the video (such as colour contrast, intensity, orientation etc.) into a saliency map, and it is from this map that predictions of gaze location are made.

Different video contexts have different gaze patterns associated with them e.g. football and open sign language [2]. Due to the omission of any top-down processes, saliency has been shown to be inadequate for such video contexts (including open sign language [10]). Other techniques [7] introduce the idea of a top-down approach, but rely instead on categorising objects within an image and then associating pre-determined probabilities of eye-fixation with each object. Although this incorporates a top-down approach it ignores prior knowledge available about the scene.

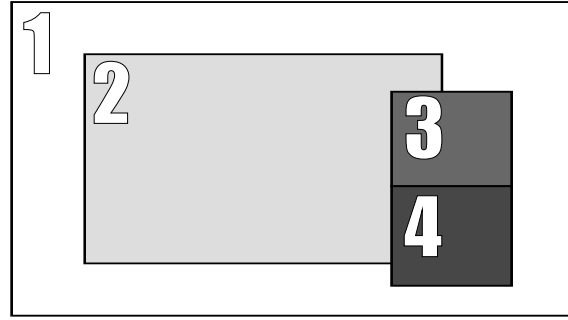
Open sign language is the term used to describe video with an in-vision signer and is used as an alternative to subtitles, which are often inadequate at conveying concepts such as emotion. Figure 1(a) shows a sample frame from open signed video footage, and figure 1(b) shows how an open signed frame can be divided up into distinct regions.

Open sign language is an excellent video context on which to perform gaze prediction due to the well-defined nature of the gaze patterns as shown in [13] [10] [1] and [3]. Previous work has been reported on analysing the gaze patterns of sign language in various forms, including videoconferencing

*Department of Electrical and Electronic Engineering, University of Bristol, Bristol, BS8 1UB, UK. Email: sam.davies@bristol.ac.uk



(a) A sample frame from one of the 'holby' open signed video sequences



(b) Diagram of the different regions in an open signed frame: 1-background; 2-programme inset; 3-signer's face; 4-signer's body

Figure 1: Open sign language sample frame and segmentation.

[3] and open sign video [10], with the aim of coding video with perceptually varying quality. Ciaramello and Hemami [8] developed an objective metric for measuring the intelligibility of sign language before using it to optimise rate distortion in coding video [9]. These studies all showed that within the sign language context it is possible to make coding gains without loss of intelligibility or perceived quality.

In order to design an open sign language gaze tracker we first perform an eye tracking study (see Section 2) which identifies several cues which might be of use to predict the location of a viewer's gaze. These cues include motion of the signer's hands, the location of faces within the frame, the orientation of the signer's head, and shot changes within the inset. We then develop methods to extract these cues from the video, including designing a grid-based likelihood ratio tracking system to predict the orientation of the signer's face. A grid-based Bayesian state estimation technique is employed to combine the cues into a probability surface which can be used to predict the viewer's gaze.

This paper is organised as follows: Section 2 describes the eye tracking study carried out with sign language material, before Section 3 investigates the suitability of possible cues for a gaze prediction routine. Sections 4 and 5 detail how two of these cues (Shot Changes and Face Orientation respectively) can be generated from the raw video, including in section 5.2 a face tracking routine based on face detections. We then use these cues to develop some gaze prediction routines in section 6, initially a single cue predictor (section 6.1) before moving onto a multicue model (section 6.2). The paper is concluded in section 7, including a discussion on further work.

2 Eye Tracking Study

Five participants, each of whom is fluent in British Sign Language (BSL), were shown a random selection of open signed clips at standard definition (720×576, 25fps). There were 30 clips in total with durations between 20 and 45 seconds, sourced from 3 separate programmes originally broadcast by the British Broadcasting Corporation (BBC), here referred to as *holby*, *meaning* and *caribbean*. The clips were recorded off-air, as MPEG-2 broadcast streams, with a bitrate of around 5 Mb/s.

Eye tracking was carried out using a tobii x50 eye tracker system, which consists of a control PC and a completely non-intrusive eye tracking device. The eye tracker uses an infra red camera to track the reflections of the pupils of the subject, and via a calibration routine the computer calculates the

gaze location. The system operates at 50Hz, so the mean of each pair of samples is taken to yield a single gaze location per frame.

The experiments were performed in a room with dimmed lighting, with the participants sitting in front of a 52 inch (16:9) plasma screen (111cm × 62cm active visible area) with the resolution set at 720×576 (Standard Definition). In order to comply with the Preferred Viewing Distance (PVD) outlined in the ITU Recommendation [18] the subject sat approximately 3.5m from the screen, as shown in 2. The eye tracker is accurate within the nearest 0.5 visual degrees, which corresponds to an accuracy on the screen of 2.98cm. This allows a gaze prediction accuracy of around 20 pixels horizontally and 28 pixels vertically.

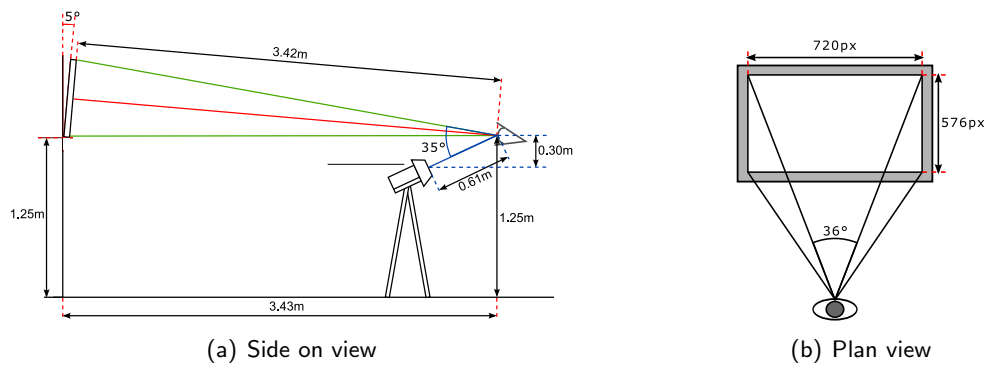


Figure 2: Eye tracking experiment set up



Figure 3: A sample frame with eye tracking fixation density plotted as hotspots

Figure 3 is a representation of the spatial distribution of the eye-tracked gaze locations and from this we see that clearly the most significant region is the signer's head. This supports previous work in both videoconferencing contexts [3] and open sign contexts [10]. The programme inset video also has gaze fixations localised on it, although figure 3 provides no information as to when these occur. On inspection of the video we notice that inset fixations tend to occur on shot changes and pauses in signing and when these saccades occur, if there are faces within the inset video then the fixations are

often localised at these positions.

3 Possible Gaze Prediction Cues

3.1 Shot Changes

When a shot change occurs in the inset it is natural for the viewer to exhibit a brief saccade from the signer to the inset in order to analyse the new scene. The true locations of the shot changes in the inset were found manually. We then investigated the proportion of these shot changes which are followed by a gaze location which isn't defined as part of the signer's face, within a given time frame.

For the *meaning* clips it was found that 68% of programme inset shot changes were followed by a saccade away from the signer's face within 25 frames (1 second).

It is then proposed that shot changes in the inset programme video will form a good cue for predicting the gaze position between the signer's face and the inset video.

3.2 Signer Signing

Conceptually it seems intuitive that when the signer is signing, the attention of the viewer will be concentrated around the signer, and conversely when there is no signing taking place the gaze will move to the inset. In order to demonstrate this hypothesis, each of the source video frames was manually classified as 'signer signing?' (yes or no).

Figure 4 shows the horizontal gaze position for an observer of a sequence (*holby_03*) and a plot of whether a signer is signing for any given frame, obtained from ground truth data. These results imply that there is a high correlation between periods of signing and gaze location (signer or programme inset).

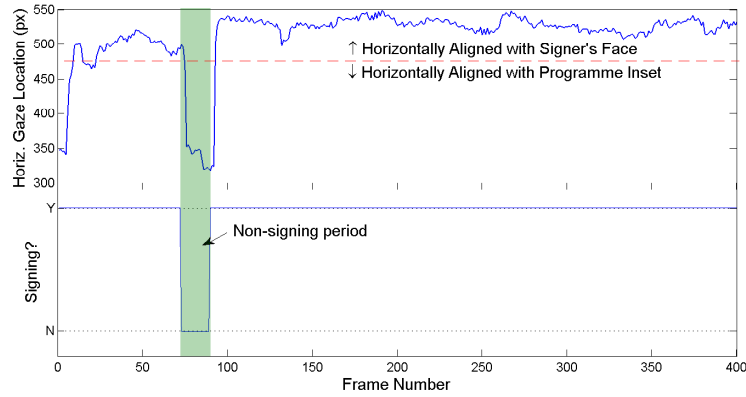


Figure 4: A comparison of horizontal gaze position with signing periods

In order to attempt to quantify this relationship between signing and gaze location we define $\mathcal{GT}_{\text{sign}}$ as the set of frames labelled as 'signing' (by manual ground truthing) and $\mathcal{ET}_{\text{sign}}$ as the set of frames in which the eye gaze locations are concentrated on the signer's face (defined as being within a pre-defined bounding box). From these we define two measurements outlined in equations 1 and 2, the former

measuring the proportion of frames where the gaze is on the signer which are labelled as signing, and the latter the proportion of the 'signing' frames which have gaze locations on the signer.

$$p_{\text{sign}}^{\mathcal{ET}} = \frac{|\mathcal{GT}_{\text{sign}} \cap \mathcal{ET}_{\text{sign}}|}{|\mathcal{ET}_{\text{sign}}|} \quad (1)$$

$$p_{\text{sign}}^{\mathcal{GT}} = \frac{|\mathcal{GT}_{\text{sign}} \cap \mathcal{ET}_{\text{sign}}|}{|\mathcal{GT}_{\text{sign}}|} \quad (2)$$

For the 7000 ground truthed frames, and 5 study participants, values of $p_{\text{sign}}^{\mathcal{ET}} = 0.95$ and $p_{\text{sign}}^{\mathcal{GT}} = 0.82$ were obtained. This implies that 95% of the time the observer is looking at the signer, signing is taking place, and that for 82% of the time signing is taking place the observers are concentrating on the signer. Therefore being able to detect whether signing is occurring for any given frame will be a good cue for predicting the gaze locality. The following possibilities were considered:-

3.2.1 Motion of the hands

The motion of the signer's hands intuitively appears to be a good cue for detecting whether or not a signer is signing at any given time.

In an attempt to measure the motion of the hands the magnitude of the h.264 motion vectors in a pre-defined region around the signer's hands is calculated. Figure 5 shows this motion metric for sequence *meaning_06*, along with the signing frames, defined from the ground truth data.

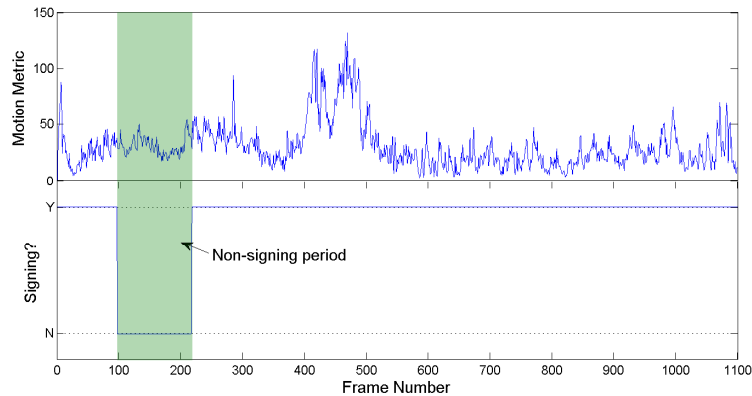


Figure 5: A comparison of signing periods with a simple metric for the motion of the signer's hands

The non-signing period is during a section of low motion compared to the global mean, and just before signing starts again there is a sudden increase in motion - these two facts could be used to help define a signing estimator from the motion in the region around the signer's hands. However, these two cues (low motion and sudden increase) are not exclusive to this non-signing region and occur elsewhere within this clip - making it difficult to localise non-signing periods.

This motion metric could be improved by segmenting out the hands of the signer and looking at the entropy of the direction of the motion (towards the end of a signing sequence, a signer will lower their hands to a stop). However initial investigations show that other, simpler cues can provide superior information.

3.2.2 Orientation of the head

Observation of open-signed material leads to the conclusion that the orientation of the signer's head might correlate well with whether or not they are signing, since when not signing the signer tends to look towards a preview monitor placed virtually to imply they are watching the inset programme. This, therefore, could in turn be a good cue for predicting the gaze location between the signer's face and the inset programme.

The source video frames were again manually classified according to 'signer's facial orientation' (frontal, profile or unknown). In over 95% of the 7000 frames classified, the face was frontal then the signer was signing. This suggests that the orientation of the signer's face is a good cue for whether or not signing is taking place in any given frame.

We also investigated whether or not the orientation of the head is a good measure for whether or not the viewers are looking at the signer. Equation 3 calculates the proportion of eye gaze locations which are correctly predicted by the facial orientation of the signer, $p_{\text{head}}^{\text{cor}}$. Here a gaze location is 'correct' if it is on the signer when the face is frontal or is not on the signer when the face is profile.

$$p_{\text{head}}^{\text{cor}} = \frac{|\mathcal{GT}_{\text{profile}} \cap \overline{\mathcal{GT}_{\text{sign}}}| + |\mathcal{GT}_{\text{frontal}} \cap \mathcal{GT}_{\text{sign}}|}{|\mathcal{ET}_{\text{sign}} \cup \overline{\mathcal{ET}_{\text{sign}}}|} \quad (3)$$

$\mathcal{GT}_{\text{frontal}}$ and $\mathcal{GT}_{\text{profile}}$ are the sets of gaze locations when the signers face is frontal and profile respectively, and $\mathcal{ET}_{\text{sign}}$ is the set of gaze locations in which the viewer is fixated on the signer (as previously).

Using the previously described ground truth data, the proportion of accurately 'predicted' locations is found to be 80%. This implies that a predictor based on the orientation of the face of the signer would yield impressive results for gaze location prediction. Therefore a model for the orientation of the face of the signer is proposed in Section 5, which uses face detections and a tracker in an attempt to discover the orientation of the head of the signer. Although based on the orientation of the signer's head, this model is ultimately attempting to predict whether a viewer will be fixating on the signer or the inset for a given frame.

4 Shot Change Detection

To detect shot changes in the inset programme video a simple technique is employed. A 32-bin histogram of the intensity of the pre-determined inset region (see figure 1(b), region 2) for each frame and the L_2 norm of the difference between consecutive frames is calculated:

$$\Delta_i = \|h_{i+1} - h_i\| = \sqrt{\sum_j (h_{i+1,j} - h_{i,j})^2} \quad (4)$$

where h_i is the histogram vector at time i , $h_{i,j}$ is the j th bin of the histogram at time i , and Δ_i is the consecutive difference between histograms at time i .

Shot changes are most likely to be related to the 2nd difference of this motion measure - i.e. a sudden large change in the difference measure (L_2) will represent a shot change. Figure 6 shows this normalised shot change detection metric (d_2) and the fixations on the inset for each of the test subjects. The 50 frame windows after each shot change cover fixations, which supports the idea that the shot changes cause fixations on the inset.

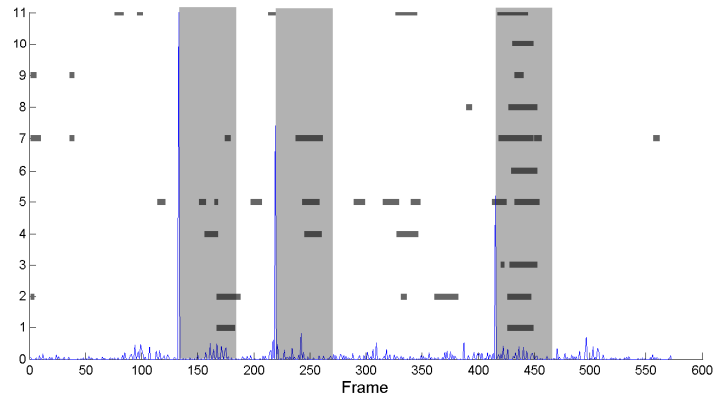


Figure 6: Shot change metric (d_2), with 50 frame windows after shot changes highlighted in light grey.

5 Face Tracker

5.1 Face Detection

The proposed gaze detector is based very heavily on face locations in a given frame for predicting the eye gaze position. We utilise the face detector proposed by Viola and Jones [19], which uses a cascade framework and Haar-like features to detect faces. The detector works on an individual frame basis, and detects faces of varying sizes based on a set of similar features. The detector is run with both frontal and profile feature cascades, so that 2 sets of detections are made. Each detection is specified by a box, within which a face has been detected. Figure 7 shows a sample frame and its associated face detections.



(a) Frame from a *'holby'* sequence, with correctly detected faces

(b) Frame from a *'meaning'* sequence, demonstrating an error with the face detection

Figure 7: Output from the face detector with the solid line representing a frontal detection, and the hashed line a profile detection

Although the detector can make correct detections (Figure 7(a)), it can also make incorrect detections, such as the profile detection on the signer's face in figure 7(b). Therefore it is not possible to simply rely on the output of the face detector. This motivates the introduction of a tracker, which will

decide how likely a face of a given orientation is at a given time.

5.2 Tracking

We use a grid based likelihood ratio tracker [11] in order to detect not only the location of a face, but also its orientation (frontal or profile) at a given time. This results in a surface which represents the likelihood ratio of a face existing at a given point versus no face being present in the frame at all.

We define states $s \in S$ as triples (s_x, s_y, s_d) , which can be thought of as multiple grid-planes, each representing a tracking direction, $s_d \in D$. Here we choose $|D|$ to be 4, representing perpendicular directions in the 2-dimensional plane (up, down, left and right). Measurements $\xi_k \in \Xi$ represent the locations of detected faces in frame k as detected by the face detector from Section 5.1, and can be expressed as a coordinate pair $(\xi_{x,k}, \xi_{y,k})$. The likelihood ratio for frame k is written $\Lambda_k(s)$, $\forall s \in S$. The likelihood ratio recursion begins with a motion update, shown in Equation 5.

$$\Lambda_k^-(s|s_d) = \mathfrak{R}(\Pi_k(s_d), s_d) * \frac{f_d}{|f_d|} \quad (5)$$

where $*$ represents a convolution, $\Lambda_k^-(s|s_d)$ is the motion update likelihood ratio for a given direction plane s_d and $\mathfrak{R}(H, v)$ rotates the plane H by $\frac{v\pi}{2}$, $v \in \{0, 1, 2, 3\}$. $\Pi_k(d)$ is the result of leaking likelihood ratio between directional planes, shown in Equation 6:

$$\Pi_k(d) = \sum_{i \in D} \Lambda_{k-1}(s|s_d = i) \cdot l(i, d) \quad (6)$$

$l(i, d)$ is the leak coefficient from direction plane i to plane d . In our implementation we choose a rotationally symmetric function such that $l(i, d) = 0.5 \cdot \delta_{\beta,0} + 0.1 \cdot \delta_{\beta,1}$, where $\delta_{i,j}$ is the kronecker delta and $\beta = |i - d| \pmod{|D|}$. $d, i \in \{0, 1, 2, 3\}$ in this work, since $|D| = 4$. f_d is a vector of filter coefficients defined in Equation 7:

$$f_d(i) = \begin{cases} 1 & 0 \leq i < c_d \\ e^{-\tau(i-c)} & c_d \leq i < n \end{cases} \quad (7)$$

$f_d(i)$ is the i th coefficient of the motion prediction filter of length n . The filter is designed such that pixel movements up to distance c_d are equally likely, after which the likelihood dies away exponentially. Values of $c_d = 3$ for horizontal direction planes, and $c_d = 1$ for the vertical planes were used, along with a decay factor of $\tau = 0.25$.

The next step of the recursion involves calculating the likelihood ratio $\mathfrak{L}_k(\xi|s_x, s_y)$ for frame k , measurement ξ , given the states defined by s_x and s_y (the measurements are independent of direction). Equation 8 shows how this is calculated.

$$\mathfrak{L}_k(\xi|s_x, s_y) = e^{-[(\xi_x - s_x)^2 + (\xi_y - s_y)^2]/2\sigma^2} \quad (8)$$

ξ is a detection at point (ξ_x, ξ_y) , and σ^2 is the variance of the 2-dimensional Gaussian.

Finally the likelihood ratio surface is updated with the new information as detailed in Equation 9.

$$\ln \Lambda_k(s) = \ln \Lambda_k^-(s) + \ln \mathfrak{L}_k(\xi_k|s) \text{ for } s \in S \quad (9)$$

This recursive system is repeated through the frames of the video and a likelihood ratio surface evolves. Since we are interested in the likelihood ratio of a detection at any point in the image plane,

we marginalise over the different directional planes to form a 2-dimensional likelihood ratio surface, $\tilde{\Lambda}_k(x, y)$, as shown in Equation 10.

$$\ln \tilde{\Lambda}_k(x, y) = \sum_{s_d \in D} \ln \Lambda_k(s|s_d) \quad (10)$$

5.3 Face Tracking Results

The face tracker is used to predict whether or not the viewers will be concentrating on the signer or elsewhere for any given frame. To this end we investigate the accuracy of G_{sign}^k . We define $G_{\text{sign}}^k = 1$ to be correct if the viewer is looking at the signer's face for frame k . From this we calculate recall, which is a measure of the proportion of false negatives recorded. In this scenario we are primarily concerned in minimising the false negatives, since it is more important to correctly predict fixations on the signer than it is to predict them on the inset. Equations 11 and 12 define precision and recall respectively, and Figure 8 shows how recall varies with d (from Equation 13) for a set of 30 sequences.

$$\text{Precision} = \frac{\text{Number Correct}}{\text{Number Correct} + \text{False Positives}} \quad (11)$$

$$\text{Recall} = \frac{\text{Number Correct}}{\text{Number Correct} + \text{False Negatives}} \quad (12)$$

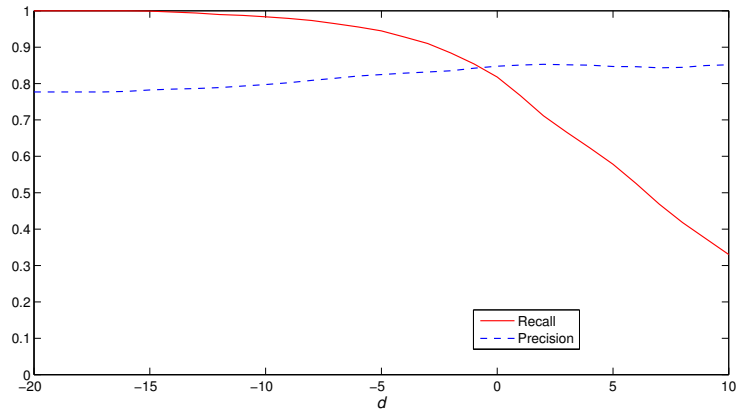


Figure 8: Average Precision and Recall for 'meaning' sequences, varying with d from Equation 13

As can be seen in Figure 8 it is possible to attain a 95% recall value with $d = -5$. This is the value of d used in this work, as it yields high recall, but a suitably small number of false positive detections (the precision is over 80%).

6 Gaze Predictor

6.1 Single Cue Gaze Predictor

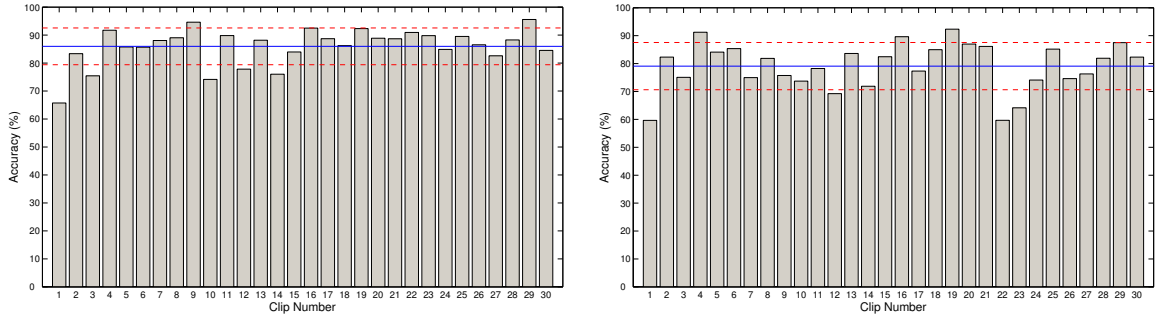
We propose a gaze predictor which is based on the face detections from Section 5.1 and the likelihood ratio tracker from Section 5.2. The tracker is used to decide whether or not the gaze is concentrated on

the signer or the inset. Two separate trackers are run on a predefined bounding box around the signer’s face simultaneously - one tracking frontal face views, the other profile. The output of each of these trackers is a likelihood surface for each frame, $\tilde{\Lambda}_k^{\text{fro}}$ for the frontal tracker and $\tilde{\Lambda}_k^{\text{pro}}$ for the profile tracker for frame k . Comparison of the maxima of these surfaces allows a prediction of whether or not the viewer will be looking at the signer or the inset. Equation 13 describes the prediction process, with d being a constant.

$$G_{\text{sign}}^k = \begin{cases} 1 & \text{if } \max(\ln \tilde{\Lambda}_k^{\text{fro}}) - \max(\ln \tilde{\Lambda}_k^{\text{pro}}) > d \\ 0 & \text{otherwise} \end{cases} \quad (13)$$

The gaze predictor specifies that if $G_{\text{sign}}^k = 1$, i.e. the signer is signing, then the gaze prediction will be at the location of the maximum of $\tilde{\Lambda}_k^{\text{fro}}$, i.e. the tracked location of the head of the signer. If $G_{\text{sign}}^k = 0$ then the viewer will be looking at the inset, and hence the predictor returns this region.

In order to assess the accuracy of the gaze predictor presented here, the results are compared with those from the eye tracking study. Since the gaze predictor returns boxes of high gaze likelihood, either the face of the signer or the inset, we say an eye track result has been correctly predicted if it falls within the predicted region for a given frame. From this we calculate a percentage of the eye tracking results that were correctly predicted by the gaze predictor.



(a) Accuracy results predicting gaze location as either signer or inset

(b) Accuracy results when localising the 'inset' predictions to face detections

Figure 9: Percentage of eye track locations correctly predicted by the gaze tracker for 30 different clips. Including the mean and a standard deviation either side of the mean.

Figure 9(a) shows the individual results for 30 different video sequences. The results vary from 70% to 95%, with the mean being 86% and the standard deviation 0.066. These results, however, only localise the gaze prediction to either the signer’s face or the inset video. In order to attempt to further localise the gaze prediction within the inset we use the face detections found in Section 5.1 - defining an eye gaze location within the inset as predicted correctly if it is within the bounding box of a face detection. Figure 9(b) shows the results for the same 30 clips as before, which vary from 60% to 95%, with the mean being 79%, and the standard deviation 0.085. As expected, the more specific the localisation the less accurate the prediction and the more varied the results.

Clips 22-24 have a significant drop in accuracy in the face-localised technique shown in figure 9(b) compared to that shown in figure 9(a). This is due to the nature of the content of the video - clips 21-30 are taken from a nature programme and therefore there are no human faces in the original video.

Therefore when the face localisation is invoked, there are no faces to find in the inset and therefore the gaze predictions are essentially random across the inset.

6.2 Multiple Cue Gaze Predictor

We use grid-based Bayesian state estimation to combine the chosen cues (detected face locations, signer's facial orientation and shot change detections). This means that each pixel has an associated likelihood of being a fixation position for any given frame, and the likelihood surface for the following frame is predicted from the current surface and the updated cues.

We work in the loglikelihood domain for ease of calculation and at each update stage ensure that the range is restricted to $[-2, 5]$. In the following explanation of the technique all of the values used here were found empirically to optimise the output of the predictor compared to eye tracking results. As a prior distribution we have a uniform base loglikelihood across the frame with a slightly higher value in the signer's bounding box and the inset (here these were 0 and 3 respectively).

The algorithm has two parts: update and prediction. The prediction stage takes the form of a 2D Gaussian filter across the entire frame - this is representing the principle that a position of high gaze likelihood for one frame will also be a likely gaze position for the following frame, but that there might be some slight variation in local position.

The update stage consists of three separate procedures, one for each of the selected cues. Initially the loglikelihood for the entire frame is reduced (here by 0.05) so that regions not highlighted as increasing the likelihood are actually reduced. Then each of the cues defines regions and increments the loglikelihood accordingly:

- **Face Locations** The face detector returns face locations and these generate regions which are of high gaze probability. We combine this cue with the signer's face orientation to weight the face locations accordingly. If the detected face appears to be the signer's face when the signer is signing (frontal face tracked) or the detected face is in the inset when the signer isn't signing (profile face tracked) then the detection is said to be 'active' and a certain loglikelihood is added to the detected regions (here chosen to be 0.25). Conversely, if a detection is labelled as 'inactive' (not signing & on signer, or signing & on inset) then a smaller loglikelihood is added to the detected facial region (here, 0.1).
- **Signer's Facial Orientation** The output of the signer's facial orientation tracker is an indication of whether or not the signer is signing at any given time (see section 5.3). We use this information to add loglikelihood uniformly across the inset when the signer isn't signing, since this event results in saccades from the signer to the inset. Figure 10 shows the form that this addition takes. k represents the point at which the signer stops signing, L_{\max} the maximum loglikelihood added to the inset, d the delay length and w the length of the window over which increases occur (the latter three were set at 0.4, 0 and 10 respectively).
- **Shot Changes** Shot changes in the inset are detected and modulate the loglikelihood of the inset in the same way as the signer's facial orientation does. Figure 10 details the way the values of the inset are increased when a shot change is detected at frame k . Here the delay (d) was set at 10, the window length (w) at 15 and the maximum increase (L_{\max}) at 0.4.

This completes the update stage and the prediction/update procedure is repeated for all subsequent frames. To generate the likelihood surface, the exponent of the value of the loglikelihood is taken for each pixel, and this new surface is normalised so that it sums to 1.

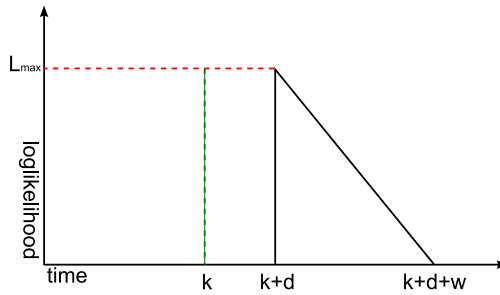
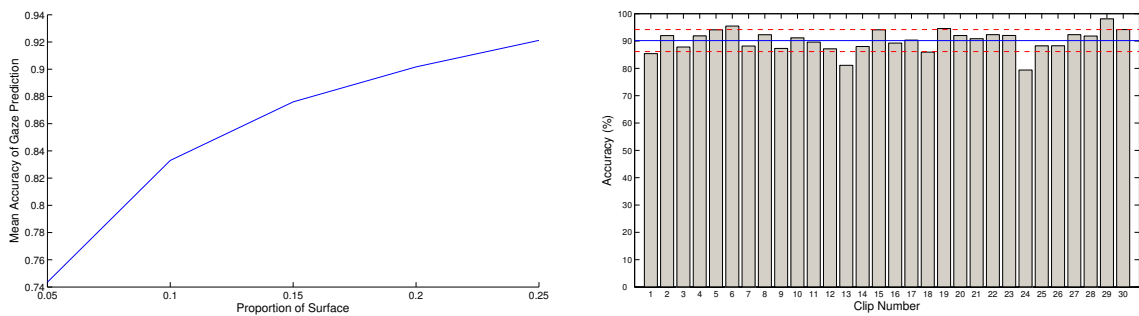


Figure 10: Loglikelihood increases over time for detected signing status changes and shot changes

Since this grid-based Bayesian approach produces a likelihood surface for each frame rather than a predicted location, direct comparison is impossible. Therefore we use a threshold-based method of estimating gaze predictions from the likelihood surface. Figure 11(a) shows the accuracy of the multi-cue gaze predictor, with a varying threshold. The threshold represents the proportion of each frame which is accepted as a gaze prediction (region of interest).

Figure 11(b) shows the proportion of correct eye gaze locations for the 30 individual clips, with the ROI threshold set at 20% of the entire frame. The mean is 90%.



(a) Accuracy of the multi-cue gaze predictor, based on a varying threshold of the proportion of each frame accepted as a prediction location

(b) Percentage of eye track locations correctly predicted by the multi-cue gaze predictor with the ROI threshold set at 20%. Includes mean and a standard deviation either side of the mean.

Figure 11: Multi-cue gaze predictor results

It can be seen from figure 11(a) that the accuracy clearly varies with the proportion of each frame labelled as a Region of Interest (ROI). For an ROI coverage of 15%, an average accuracy of 87% is obtained, whilst with a coverage of 20% this accuracy is increased to 90%, with a standard deviation of 0.044. In comparison, a prediction model which estimates that the viewer will be continuously fixated on the signer's face has an average accuracy of 73%, whilst the single-cue techniques yield an 87% accuracy for inset-vs-signer (see figure 9(a)) and an 81% accuracy for localised faces (see figure 9(b)), where each technique has an ROI coverage area of 15-20%. Therefore the multi-cue approach yields not only more accurate results than both the single-cue techniques and the static prediction, but also reduces the variance of the prediction accuracy.

The multi-cue technique therefore not only outperforms the single-cue gaze predictor but also pro-

vides more information in the form of a likelihood surface as opposed to discrete gaze prediction locations. This surface provides a relative importance score across each frame - not just a binary ROI/nROI map, and this surface can be incorporated into the same end uses as the binary map can (e.g. variable quality coding).

7 Conclusion

In this paper we have proposed a gaze prediction system for a specific video context, namely open sign language. This system uses a grid-based Bayesian state estimator to combine various cues, which themselves have been generated from the raw video. We have also investigated the validity of different cues (shot change in the inset programme video, location of faces across the frame, the orientation of the signer's face and the motion of the signer's hands), and developed techniques for extracting these cues from the video at a signal level, including a novel use of a likelihood ratio tracker to detect the orientation of a face.

The proposed gaze prediction model produces a probability surface for each frame - showing the likelihood of a viewer looking at any point within the frame. The output of this was analysed against eye tracking data for a large selection of clips and was shown to outperform both a uniform probability surface, and a surface generated using face locations within the frame.

This probability surface can be used for various different purposes within coding, including bit-rate allocation creating a perceptually optimised video codec, and error protection. This gaze prediction framework can be adapted to suit different video contexts using prior knowledge of the gaze pattern.

References

- [1] D. Agrafiotis, N. Canagarajah, D.R. Bull, and M. Dye. Perceptually optimised sign language video coding based on eye tracking analysis. *Electronics Letters*, 39(24):1703–1705, 27 Nov. 2003.
- [2] D. Agrafiotis, S. J. C. Davies, N. Canagarajah, and D. R. Bull. Towards efficient context-specific video coding based on gaze-tracking analysis. *ACM Trans. Multimedia Comput. Commun. Appl.*, 3(4):1–15, 2007.
- [3] D. Argrafiotis, N. Canagarajah, D.R. Bull, H. Twyford, J. Kyle, and J. Chung How. Optimised sign language video coding based on eye-tracking analysis. *Visual Communications and Image Processing, Proc. of SPIE*, 5150:1244–1252, 2003.
- [4] E. Barth, J. Drewes, and T. Martinez. Dynamic predictions of tracked gaze. *Signal Processing and Its Applications, Proc. 7th Int'l Symposium on*, 1:245–248, July 2003.
- [5] G. Boccignone and M. Ferraro. Modelling gaze shift as a constrained random walk. *Physica A*, 331:207–218, 2004.
- [6] W. Cheng, W. Chu, and J. Wu. A visual attention based region-of-interest determination framework for video sequences. *IEICE Transactions on Information and Systems*, E88-D(7):1578–1586, 2005.
- [7] D.A. Chernyak and L.W. Stark. Top-down guided eye movements. *IEEE Trans. on Systems, Man and Cybernetics, Part B*, 31(4):541–522, Aug 2001.

- [8] Francis M. Ciaramello and Sheila S. Hemami. "can you see me now?" an objective metric for predicting intelligibility of compressed american sign language video. volume 6492, page 64920M. SPIE, 2007.
- [9] Frank M. Ciaramello and Sheila S. Hemami. Complexity constrained rate-distortion optimization of sign language video using an objective intelligibility metric. *Visual Communications and Image Processing 2008*, 6822(1):682213, 2008.
- [10] S. J. C. Davies, D. Agrafiotis, C. N. Canagarajah, and D.R. Bull. Perceptually optimised coding of open signed video based on gaze prediction. *Electronics Letters*, 39(21):1135–1136, October 2007.
- [11] David L. Hall and James Llinas, editors. *Handbook of Multisensor Data Fusion*. CRC Press, 2001.
- [12] L. Itti. Automatic foveation for video compression using a neurobiological model of visual attention. *IEEE Trans. on Image Processing*, 13(10):1304–1318, Oct 2004.
- [13] Laura J. Muir and Iain E. G. Richardson. Perception of sign language and its application to visual communications for deaf people. *J. Deaf Stud. Deaf Educ.*, 10(4):390–401, 2005.
- [14] E. Niebur and C. Koch. *The Attentive Brain*, chapter Computational Architectures for Attention, pages 163–186. MIT Press, Cambridge, Mass., 1998.
- [15] C. Privitera and L.W. Stark. Algorithms for defining visual region-of- interest: Comparison with eye fixations. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(9):970–982, September 2000.
- [16] Claudio M. Privitera and Lawrence W. Stark. *Neurobiology of Attention*, chapter Scanpath Theory, Attention, and Image Processing Algorithms for Predicting Human Eye Fixations, pages 24–28. Elsevier, 2005.
- [17] Umesh Rajashekar, Ian van der Linde, Alan C. Bovik, and Lawrence K. Cormack. Foveated analysis and selection of visual fixation in natural scenes. In *IEEE International Conference on Image Processing*, pages 453–456, 2006.
- [18] International Telecommunications Union. Recommendation itu-r bt.500-11: Methodology for the subjective assessment of the quality of television pictures, 2002.
- [19] Paul Viola and Michael J. Jones. Robust real-time face detection. *International Journal of Computer Vision*, 57(2):137–154, May 2004.