IJCNN 2019. International Joint Conference on Neural Networks. Budapest, Hungary. 14-19 July 2019

Your Eyes Say You're Lying: An Eye Movement Pattern Analysis for Face Familiarity and Deceptive Cognition

Jiaxu Zuo[1], Tom Gedeon*[1], and Zhenyue Qin[1]

[1]Australian National University, Canberra, Australia

Abstract

Eye movement patterns reflect human latent internal cognitive activities. We aim to discover eye movement patterns during face recognition under different conditions of information concealment. These conditions include the degrees of face familiarity and deception or not, namely telling the truth when observing familiar and unfamiliar faces, and deceiving in front of familiar and unfamiliar faces. We apply Hidden Markov models with Gaussian emission to generalise regions and trajectories of eye fixation points under the above four conditions. Our results show that both eye movement patterns and eye gaze regions become significantly different during deception compared with truth-telling. We show the feasibility of detecting deception and further cognitive activity classification using eye movement patterns.

1. Introduction

During criminal or other forms of investigations for justice, suspects may deceive the investigators by claiming to not recognise a familiar face in order to remove them from suspicion or to protect their fellow co-conspirators. Failure to detect this deception can cause a threat to society since it can lead to imprisoning innocent citizens whilst let the guilty go free [21]. For example, several days before the terror attack in Brussels in 2016, one of the accomplices was arrested and interrogated about his relationship with the terrorist group. He denied any familiarity with the photos of that terrorist group shown to him [13].

Lie detection, such as to interrogate suspects about the familiarity of other people, plays an essential role in maintaining justice and stability of society. However, human capability to discern deception is poor, with an accuracy of slightly better than chance [3]. Thus, people have proposed various methodologies aiming to support criminal investigators for fighting such tricks. It has been shown that lying will lead to a range of physiological changes of the body [1]. These changes caused by telling lies facilitate the emergence of physiological approaches to lie detection [20]. For example, functional Magnetic Resonance Imaging (fMRI) based methods can present very accurate results [5]. Nonetheless, these methods contain too much noise, high expense, and have other drawbacks which make them infeasible to use in practice [7].

Another perspective of discerning deception is through leveraging behavioural cues, which may seem to be negligibly perceived by normal people [17]. For example, facial micro-expressions, like eyebrows raising, may reflect that the subject is trying to hide his or her real emotions [6]. Nevertheless, it is very laborious to train experts to read these behavioural skills due to the significant variability among different subjects [21]. That is, the same micro-emotion may have different cognitive meanings with respect to different persons.

Unlike fMRI and facial micro-expressions, eye fixation locations and trajectories are relatively easy to capture and collect by ordinary people with the assistance of simple eye trackers. Previous studies indicate that trajectories of eye gaze movements reflect individuals' underlying cognitive activities [9]. For instance, previous research shows that there will be fewer eye fixations during the (normal) process of recognising familiar faces than unfamiliar faces [8].

Another key feature of eye movements lies in its involuntary property. That is, participants are not able to easily alter their behaviours, irrespective to the undertaken tasks [18]. Due to these advantages of utilising eye movements, we study the variability of eye movements for deception and truth telling given a familiar or unfamiliar face. This variability includes the visiting orders of facial regions and the distributions of eye fixation locations.

Several recent publications have attempted to investigate the relationship between face familiarity and eye fixations. Millen et al. collected the number of fixations in total and in different facial regions of stating to be unfamiliar under four

*T. Gedeon (tom@cs.anu.edu.au) is the corresponding author.

978-1-7281-2009-6/$31.00 paper N-19623.pdf

conditions, namely unfamiliar faces, newly learned faces, famous celebrity faces, and personally familiar faces [14]. However, they did not include the analysis of eye movement paths and the facial regions of interest in their work. Lancry-Dayan et al. recently explored the possibility of detecting a subject's familiarity with other faces using machine learning techniques by demonstrating his or her four face photos in parallel [12]. They discovered that among the given four faces, subjects tend to fixate more on familiar faces, followed by a tendency to move eyes away from it [12]. Furthermore, machine learning techniques can successfully leverage this behaviour and showed robust results among different individuals [12]. Nevertheless, they did not explicitly scrutinise the eye movement behaviours on a single individual's face. That is, their study was to investigate the eye movement patterns on four parallel displayed photos instead of single photos.

In this paper, we aim to investigate the eye movement behaviors during different face recognition tasks under distinct cognitive tasks, namely truth telling or lying. Specifically, we explore the eye gaze patterns under four different situations, which are: (a) telling truth on a familiar face; (b) telling truth on an unfamiliar face; (c) lying on a familiar face; and (d) lying on an unfamiliar case, which is unlikely in the real world. Unlike the previously published work, we focus on the trajectories of eye fixations and their distributions on facial regions of interest when undertaking different cognitive tasks. That is, we show that individuals stare at different facial regions when telling the truth or a lie (particularly on familiar faces), and the visiting orders of these regions are different.

Our work can contribute to the future design of automatic deception detection systems. For example, Wu et al. take inputs of multiple modalities, including individual motions, audio and so on, to discern the veracity of expressions from real-life courtroom trial videos [21]. Our work may serve as an additional input channel. Furthermore, our study may also shed light into the field of visual saliency predictions within computer vision, which focuses on predicting which objects that people will fixate on, given an image or video [11]. Our results indicate that additional channels of people's cognitive activities may further increase the saliency prediction accuracy. Moreover, our research also reveals the possibility of interpreting people's cognitive activities from eye movement patterns.

2. Methods

2.1. Fixation Identification

The raw eye gaze data is in the form of the Cartesian coordinate system, namely its locations are expressed by (x, y) coordinates. It is essential to separate and label eye tracking points as fixations or saccades, since improper classification can have a dramatic influence on higher-level analyses [19].

The dispersion threshold algorithm is considered to be the most robust and accurate approach for identifying fixations and saccades in eye-tracking protocols compared with other approaches such as velocity-based and area-based ones [19]. One essential parameter for this algorithm is called the pixel tolerance, which is a threshold that determines whether to classify a new eye-gaze point as a fixation or a saccade. Previous research reported that a normal fixation duration is about 200-250 ms [15], and each of our stimulus videos lasts for five seconds. Thus, we can deduce that 20 fixation points per video should be a reasonable number, which leads to a setting of the pixel tolerance being 5.

2.2. Eye Movement Analysis

We apply probabilistic graphical models, specifically Gaussian Hidden Markov models (HMMs with 2D Gaussian emission distribution), to train the previous collected eye movement data from participants. This technique is broadly utilised to model data generated from Markov processes [2, 4]. That is, the next state of a process only depends on its previous state. For every pair from a current state S to its subsequent state S' in the next time step, or to the observable state in the same time step O, there is an associated probability indicating the likelihood of transforming S to S' or S to O, which can be formalised as a matrix. Finally, a vector of prior probabilities reveals the different probabilities of being in distinct starting states. We hypothesise that we will obtain different representative eye movement patterns during face recognition with distinct latent mental activities, namely telling the truth or lying. In this context, states represent different ROIs of a face image, observable states are fixation points, and the transition from the current latent state to the next state is a saccade. This is a quite direct model of eye gaze.

This methodology requires determining the number of facial regions of interest (ROIs) beforehand. We pick the number as three, since previous studies show that there are three information-rich inner regions of the face (eyes, nose, and mouth), which are particularly essential for person recognition [16].

In our case, we set the concentration parameters of Dirichlet prior distributions of initial distribution and transition probability to be 0.01. The prior covariance was set as an isotropic covariance matrix with standard deviation of 14 to ensure that the ROI has roughly the same size as the facial features.

We first trained a Gaussian-HMM using the fixation data from one participant who was viewing images under a certain cognitive condition. For example, the participant has been instructed to tell the truth, and then showed some fa-

Figure 1: One original face image and its processed result.

miliar faces. Fixation sequences during every observation is considered as viewing sequences. Followed by applying the Gaussian-HMM algorithm on multiple observation sequences, we can get a representative Gaussian-HMM with three states. Each state consists of three means and three covariance matrices representing the ROI centres and ranges.

For every participant, there are four Gaussian-HMMs representing four different eye movement strategies, corresponding to the two degrees of face familiarity and lying or not. As a result, we have 84 Gaussian-HMMs from 21 participants, and classify them into four classes based on the face familiarity and lying conditions. Afterwards, we train another Gaussian-HMM for each category in order to find a class representative Gaussian-HMM and clustering centroids of its ROIs. Eventually, we will have four representative Gaussian-HMMs, each will contain its clustering centroids of ROIs, shapes of ROIs, an initial distribution vector of ROIs, a transition probability matrix, and a Viterbi path.

3. Experiments

3.1. Face Image Preprocessing

The variations of brightness and colours of photos can cause differences in human eye saccades and fixations, as well as a significant latency of the reaction time for viewing the same images [10]. We wish to minimise the noises resulting from this process, known as an adaptation in ocular physiology. Therefore, we normalise all the experimental face images to be consistent with the same grey-scale and size of 447×335. Furthermore, we apply an oval mask to all the other parts of the images, including hair, background and so on, except the main face area. This is to avoid the eye gaze being distracted by the background and objects other than our interest regions, namely human faces. One original face image and its processed result is shown in Figure 1.

3.2. Face Images and Participants

We collected a total of 20 different grey-scale frontal-view face images, each will be displayed two times in a random order, one corresponding to telling the truth and the other for lying. That is, one participant will observe 40 images out of 20 correspond to university lectures from two distinct schools. Correspondingly, participants were sampled from these two Schools to ensure students from a particular faculty can recognise about half of the lecturers' faces, those from their own school, and have little or no knowledge for the other half, from the other school. The remaining five images were sourced from an online face database and so guaranteed to be unknown to all participants. In total, we recruited 21 participants.

3.3. Experimental Procedure

We developed a web-browser based user interface to display our collected face images to the participants. During the experimental procedure, a participant will firstly see a huge instruction indicating either to "tell the truth!" or "lie!" for the upcoming face image. For example, if the face to display is familiar to the participant, he or she is expected to pretend to have no knowledge about the shown image. The same instruction will be shown again to the participant before he or she gives the answer in order to avoid unexpected honesty. We record participants' eye fixations trajectories during this process.

4. Results

The average fixation duration of the participants was 263.2 ms (SD = 49.3 ms). That is, the average number of fixation points in one trial of the experiment is 19. We analysed a general eye movement pattern that summarises all the 84 individual HMMs and the four complete experimental settings of different combinations face familiarities and information concealing degrees, namely

1. Telling the truth given a familiar face
2. Telling the truth given an unfamiliar face
3. Lying given a familiar face
4. Lying given an unfamiliar face

We next show the detailed eye movement patterns under different experimental settings, and show that eye gaze trajectories and ROIs differ given different cognitive conditions during face recognition, namely different degrees of face familiarity and information concealing minds. To be more intuitive, all the left and right directions below are from the viewer's perspective.

4.1. General Eye Movement Pattern

Following the methods given in 2.2, we obtained 84 Gaussian-HMMs, each with three means and covariance matrices to parametrize its three ROIs. We utilised the VHEM to group these 84 HMMs into a single representation that summarises all the individual HMMs. Figure 2 and

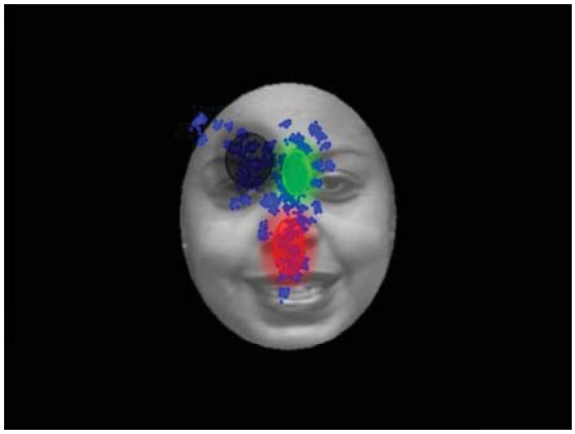

Figure 2: The general representative HMM that summarises all the 84 individual HMMs. Blue dots are fixation points. Different colours represent distinct ROIs.

Prior values	Red	Green	Blue
	0.0000	0.8475	0.1525
Transition probabilities	To Red	To Green	To Black
From Red	0.9704	0.0086	0.0210
From Green	0.0449	0.9248	0.0303
From Blue	0.0411	0.0248	0.9340

Table 1: Transition probabilities of the general representative HMM that summarises all the 84 individual HMMs.

Table 1 show the representative HMM that summarises all the 84 participants' HMMs.

From Figure 2 and Table 1, the general scan path during this face recognition mostly likely starts from the green ROI (an oval region close to the inner canthus of the right eye), then goes to the red region (an oval region covers tip of the nose and the philtrum). Afterwards, the next fixation tends to remain in the red region. Also, the path is possible to move to the black region (close to the left eye, which is an oval region slightly above the eye and covers part of the left eyebrow). Finally, the probability of returning to the initial region (i.e. the green region) is quite low. In contrast, if starting the scan path from the black region with a low probability, it is highly possible that the path remains in the black region.

4.2. Telling the Truth Given a Familiar Face

As before, we applied VHEM and generated the HMM for the case of participants telling the truth given familiar faces. Figure 3, 4 and Table 2 illustrate more details on ROI distributions and scan moving paths. Specifically, the ROI distributions on the face of this case are:

Prior values	Red	Green	Blue
	0.8626	0.0479	0.0895
Transition probabilities	To Red	To Green	To Black
From Red	0.8859	0.0402	0.0738
From Green	0.0285	0.9444	0.0272
From Blue	0.0903	0.0252	0.8845

Table 2: Transition probabilities of the ROIs for eye movement pattern during telling the truth given a familiar face.

- Red Region: Close to the inner canthus of the left eye from the viewer perspective, which is slightly above the eye, covering the inner part of the left eye and the left eyebrow, and the left upper eyelid. The semi-major axis is on the horizontal direction, while the semi-minor axis is on the vertical direction.
- Black Region: Close to the inner canthus of the right eye, which is slightly above the eye, covering a small inner part of the right eye, the right eyebrow and the right upper eyelid. The semi-major axis is on the vertical direction, while the semi-minor axis is on the horizontal direction.
- Green Region: The tip of the nose, the philtrum and most parts of the mouth. The mean axis of the oval region is slightly to the left compared to the middle axis of the face. The semi-major axis is on the vertical direction, while the semi-minor axis is the on horizontal direction.

In this sub-group, the first fixation tends to be in the red region (close to the left), followed by a high probability to remain in the red region, or it can move to the black region. Then it is highly probable to stay in the black region, if not, it can either move back and forth between the initial red region, or it can move to the green region, although the probability of moving to the green region is quite low. If the scan path starts from the black region by a low chance, the next fixation can possibly stay in the same black region, or it can move back and forth between the red and black regions. In this case, the probability of moving to the green region is still low. If the eye movement starts from the green region by a low probability, it is very likely to stay in the same region. The generated Viterbi path starts from the red region, then the black region, and concludes in the black region.

4.3. Telling the Truth on an Unfamiliar Face

We applied the same approach to generate HMMs from participants telling the truth in front of unfamiliar faces. In this case, the specific ROI distributions on the face are:

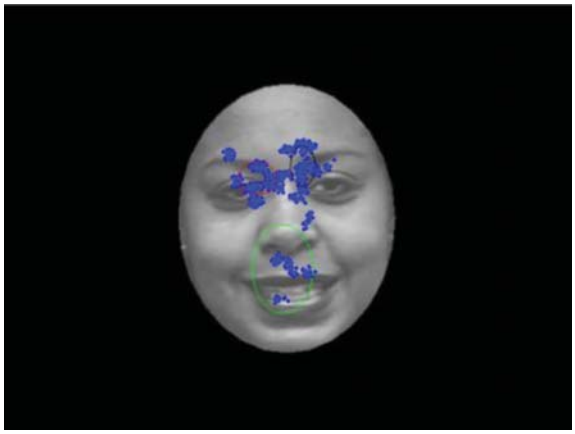

Figure 3: The ROIs for the eye movement pattern during telling the truth given a familiar face.

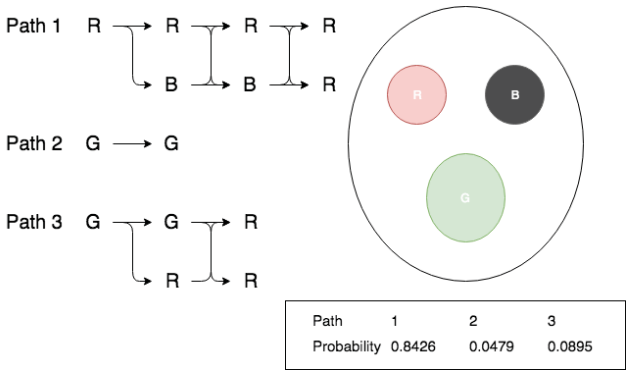

Figure 4: Scan paths for eye movement pattern during telling the truth given a familiar face.

- Red Region: A small part of the nose tip, a small part of the philtrum, and a small part of the upper lip in the middle. The mean axis of this oval region is generally in the middle axis of the face. The semi-major axis is on the vertical direction, while the semi-minor axis is on horizontal direction.
- Black Region: Close to the inner canthus of the right eye, which is being in a more inferior position to the right inner canthus, covering a small region close to the right inner canthus. The semi-major axis is at the vertical direction, while the semi-minor axis is on the horizontal direction.
- Green Region: A right part of the left eye, which is slightly to the central axis of the face, covering the most inner part of the left eye, the left eyebrow, and the region close to the left eye including a part of the nose bridge. The region is an approximate circle.

The most probable scan path starts from the black region (close to the right eye), with a high probability of remaining in the black region, or it can move to the red region or the green region (very unlikely). If it moves to the red region, next most likely action tends to stay in the same ROI, or it has a low probability to move to the green region. Otherwise, if it moves to the green region, it will likely to remain in the same region, to move to the red region by low chance. The probability of moving back to the black region is very low. In addition, it can start from green region, followed by a high probability of remaining in the green region. The probability of starting from the red region is very low, compared with from the other two ROIs. The Viterbi Path in this case implies that the most likely path is similar as the analysis above, of which starts from the black region, then comes to the red region, and end up with the green region.

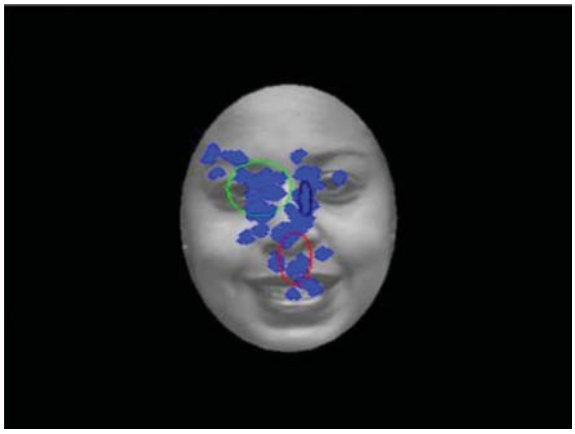

Figure 5: The ROIs for eye movement pattern during telling the truth given an unfamiliar face.

Prior values	Red	Green	Blue
	0.0820	0.3354	0.5826
Transition probabilities	To Red	To Green	To Black
From Red	0.9680	0.0215	0.0105
From Green	0.0518	0.9225	0.0257
From Blue	0.0420	0.0898	0.8682

Table 3: Transition probabilities of the ROIs for eye movement pattern during telling the truth given an unfamiliar face.

4.4. Lying Given a Familiar Face

This situation shows a significant difference compared with the previous two cases. Figure 7 illustrates the representative Gaussian-HMM when participants lie about the familiarity while viewing the familiar face. As can be seen, the ROI distributions in face recognition of this case are:

- Red Region: A small part of the nose bridge between

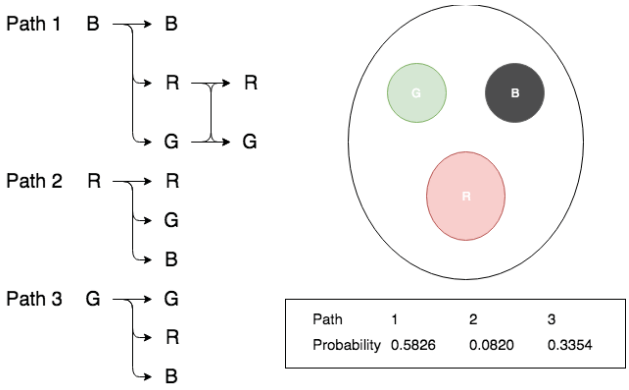

Figure 6: Scan paths for eye movement pattern during telling the truth given an unfamiliar face.

two eyes. The mean axis of this oval region is slightly to the right of the middle face axis. The oval region is an approximate circle.

- Black Region: The region between two eyebrows, covering part of the forehead close to the central of two eyebrows, which is slightly to the left of the middle axis of the face. The semi-major axis is in the horizontal direction, while the semi-minor axis is in the vertical direction.
- Green Region: Covering the lower part of the nose bridge, the tip of the nose, and the philtrum. The mean axis of this oval region is generally of the middle axis of the face. The semi-major axis is on the vertical direction, while the semi-minor axis is on the horizontal direction.

In this sub-group, the scan path most likely starts from the red region. Afterwards, it is more likely to remain in the same ROI, or with a lower probability, to move to the green or black regions, and stay the same locations subsequently. It is also possible to start from the black region, then with a higher likelihood of remaining the same region, or move to the other two regions. This Viterbi Path generally starts from the red region, then moves to the green region and converges to it.

4.5. Lying Given an Unfamiliar Face

This case also displays quite different eye behaviour compared with lying on a familiar face. As Figure 10 illustrates, the eye fixation points mainly locate on two eye regions, unlikely previously that there are a large proportion of points distribute to the nose region. The detailed ROI distributions are:

- Red Region: A small area located at the top-right left eye, slightly touching the left eyebrow, which is an approximate circle.

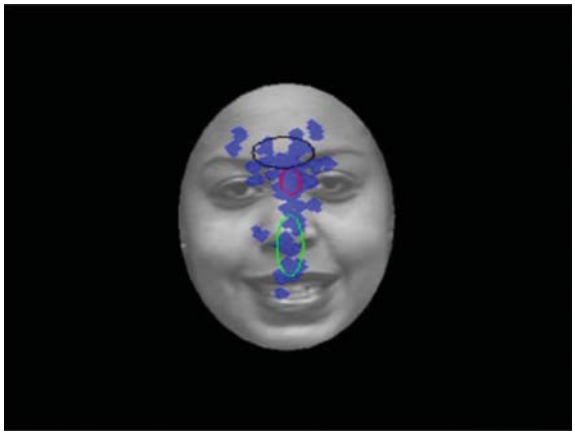

Figure 7: The ROIs for eye movement pattern during lying given a familiar face.

Prior values	Red	Green	Blue
	0.6456	0.0000	0.3544
Transition probabilities	To Red	To Green	To Black
From Red	0.9102	0.0500	0.0398
From Green	0.0119	0.9663	0.0218
From Blue	0.0305	0.0501	0.9194

Table 4: Transition probabilities of the ROIs for eye movement pattern during lying given a familiar face.

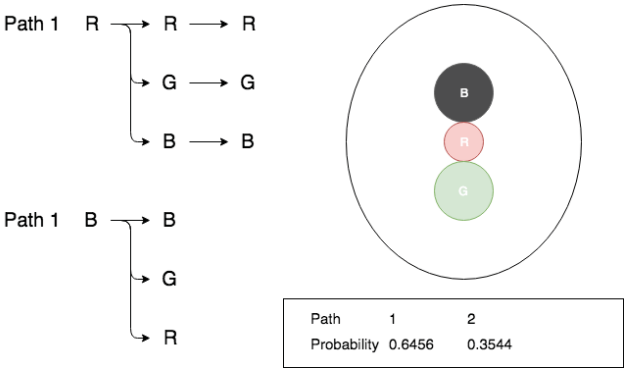

Figure 8: Scan paths for eye movement pattern during lying given a familiar face.

- Black Region: Covering a large area in the middle of the face including the nose, the philtrum, the lips, and the proximity regions. However, the fixation points are relatively sparse compared with the other two ROIs.
- Green Region: A approximate circle shape distributed at the left part of the right eye, below the right eyebrow.

In this case, it is equally likely for eye fixation points to

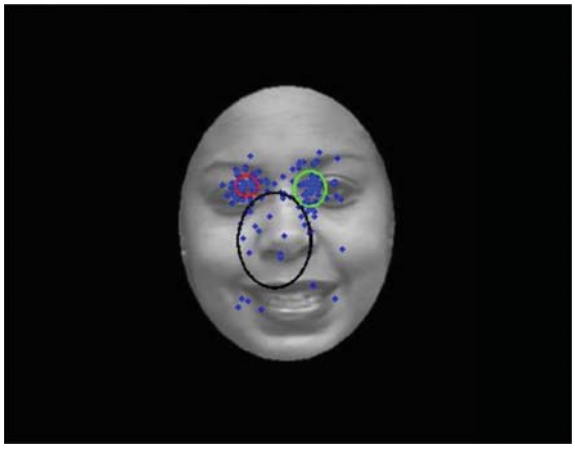

Figure 9: The ROIs for eye movement pattern during lying given an unfamiliar face.

Prior values	Red	Green	Blue
	0.4848	0.4027	0.1125
Transition probabilities	To Red	To Green	To Black
From Red	0.5602	0.3775	0.0623
From Green	0.3153	0.5622	0.1225
From Blue	0.2099	0.2751	0.5150

Table 5: Transition probabilities of the ROIs for eye movement pattern during lying given an unfamiliar face.

locate in the red or green regions. If starting from the red region, the scan path is highly likely to remain in the same region, or move to the green region, move back and forth. Moving to the black region is relatively unlikely to happen. The Viterbi path shows that this scan path is probable to conclude with an oscillation between the red and green regions.

4.6. Model Generalisation Test

In this section, we show our tested results for the generalisation of our proposed eye movement patterns for different face familiarity and information concealing. We generated a Viterbi path through the HMM methodology for each of our 21 participants. This is followed by calculating four Euclidean distances between this Viterbi path and the general Viterbi paths corresponding the four conditions, namely truth telling or lying in front of familiar and unfamiliar faces.

We evaluated the accuracy results of classifying eye movement patterns into telling the truth or a lie for both familiar and unfamiliar face recognition. The outcomes are well-beyond the chance, as Tables 7 and 8 demonstrate.

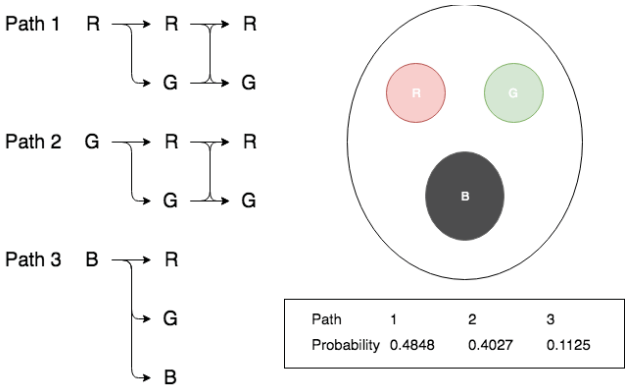

Figure 10: Scan paths for eye movement pattern during lying given an unfamiliar face.

	Familiar & Truth	Unfamiliar & Truth	Familiar & Lie	Unfamiliar & Lie
Accuracy	66.67% (14/21)	61.90% (13/21)	76.19% (13/21)	71.43% (15/21)

Table 6: Accuracy results of classifying eye movement patterns to their source face familiarity and degrees of information concealing.

	Familiar & Truth	Familiar & Lie
Accuracy	80.95% (17/21)	85.71% (18/21)

Table 7: Accuracy results of classifying eye movement patterns of familiar face recognition when telling the truth or a lie.

	Unfamiliar & Truth	Unfamiliar & Lie
Accuracy	71.43% (15/21)	80.95% (17/21)

Table 8: Accuracy results of classifying eye movement patterns of unfamiliar face recognition when telling the truth or a lie.

Furthermore, we also classify it to one of the four conditions according to the shortest Euclidean path. The 1-in-4 classifying accuracies are in Table 6. Again, as this table shows, eye movement patterns are more distinguishable when participants are lying. These results reveal a promising possibility of utilising eye movement patterns to analyse latent cognitive activities.

5. Conclusion

In this study, we generalise representative eye movement patterns, characterised by their trajectories and fixa-

tion point distributions. These patterns include telling the truth or a lie in front of a familiar or unfamiliar face, and a general one using all the individual data. We found that the eye movement patterns of lying, during both familiar and unfamiliar face recognition, are significantly different compared with truth-telling situations. Subsequent tests using the four eye movement patterns generated using HMMs with Gaussian emission, demonstrated good performance for discerning deception.

In the future, we will compose our eye movement patterns into deception detection systems as additional channels to see whether they can help to enhance the original performance of these systems. Moreover, we will test the feasibility of whether eye movement patterns behave differently under different cognitive modes such as happiness and anger.

References

[1] M. R. Bhutta, M. J. Hong, Y.-H. Kim, and K.-S. Hong. Single-trial lie detection using a combined fnirs-polygraph system. *Frontiers in Psychology*, 6:709, 2015.

[2] C. M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg, 2006.

[3] C. F. Bond Jr and B. M. DePaulo. Accuracy of deception judgments. *Personality and social psychology Review*, 10(3):214–234, 2006.

[4] T. Chuk, A. B. Chan, and J. H. Hsiao. Understanding eye movements in face recognition using hidden markov models. *Journal of vision*, 14(11):8–8, 2014.

[5] C. Davatzikos, K. Ruparel, Y. Fan, D. Shen, M. Acharyya, J. Loughead, R. Gur, and D. D. Langleben. Classifying spatial patterns of brain activity with machine learning methods: application to lie detection. *Neuroimage*, 28(3):663–668, 2005.

[6] P. Ekman. *Telling lies: Clues to deceit in the marketplace, politics, and marriage (revised edition)*. WW Norton & Company, 2009.

[7] M. J. Farah, J. B. Hutchinson, E. A. Phelps, and A. D. Wagner. Functional mri-based lie detection: scientific and societal challenges. *Nature Reviews Neuroscience*, 15(2):123, 2014.

[8] D. E. Hannula, R. R. Althoff, D. E. Warren, L. Riggs, N. J. Cohen, and J. D. Ryan. Worth a glance: using eye movements to investigate the cognitive neuroscience of memory. *Frontiers in human neuroscience*, 4:166, 2010.

[9] M. A. Just and P. A. Carpenter. Eye fixations and cognitive processes. *Cognitive psychology*, 8(4):441–480, 1976.

[10] M. Kalloniatis and C. Luu. Light and dark adaptation. 2007.

[11] S. S. Kruthiventi, V. Gudisa, J. H. Dholakiya, and R. Venkatesh Babu. Saliency unified: A deep architecture for simultaneous eye fixation prediction and salient object segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5781–5790, 2016.

[12] O. C. Lancry-Dayan, T. Nahari, G. Ben-Shakhar, and Y. Pertzov. Do you know him? gaze dynamics toward familiar faces on a concealed information test. *Journal of Applied Research in Memory and Cognition*, 2018.

[13] J. Lichfield. Paris attacker salah abdeslam denied knowledge of brussels prior to belgium terror strikes and 'wasn't asked about future attacks', 2016. [Online, accessed 2018-10-28].

[14] A. E. Millen, L. Hope, A. P. Hillstrom, and A. Vrij. Tracking the truth: the effect of face familiarity on eye fixations during deception. *The Quarterly Journal of Experimental Psychology*, 70(5):930–943, 2017.

[15] M. Miyao, S. S. Hacisalihzade, J. S. Allen, and L. W. Stark. Effects of vdt resolution on visual fatigue and readability: an eye movement approach. *Ergonomics*, 32(6):603–614, 1989.

[16] C. O'Donnell and V. Bruce. Familiarisation with faces selectively enhances sensitivity to changes made to the eyes. *Perception*, 30(6):755–764, 2001.

[17] M. O'Sullivan and P. Ekman. 12 the wizards of deception detection. *The detection of deception in forensic contexts*, page 269, 2004.

[18] J. D. Ryan, D. E. Hannula, and N. J. Cohen. The obligatory effects of memory on eye movements. *Memory*, 15(5):508–525, 2007.

[19] D. D. Salvucci and J. H. Goldberg. Identifying fixations and saccades in eye-tracking protocols. In *Proceedings of the 2000 symposium on Eye tracking research & applications*, pages 71–78. ACM, 2000.

[20] J. J. Walczyk, F. D. Igou, L. P. Dixon, and T. Tcholakian. Advancing lie detection by inducing cognitive load on liars: a review of relevant theories and techniques guided by lessons from polygraph-based approaches. *Frontiers in psychology*, 4:14, 2013.

[21] Z. Wu, B. Singh, L. S. Davis, and V. Subrahmanian. Deception detection in videos. In *AAAI Conference on Artificial Intelligence (AAAI)*, Hilton New Orleans Riverside, New Orleans, USA, 2018. AAAI Press.