Lie Detection in Virtual Reality

Judsen Hembree
Clemson University
Clemson, SC, USA
jhembre@g.clemson.edu

Jason Strickland
Clemson University
Clemson, SC, USA
jstric8@g.clemson.edu

Kevin Lin
Clemson University
Clemson, SC, USA
klin2@g.clemson.edu

Zachary Norman
Clemson University
Clemson, SC, USA
znorman@g.clemson.edu

ACM Reference Format:
Judsen Hembree, Jason Strickland, Kevin Lin, and Zachary Norman. 2022. Lie Detection in Virtual Reality. In *Proceedings of ACM Conference (Conference'17).* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

ABSTRACT

This paper analyzes the validity of lie detection in VR using the Vive Pro Eye HMD. A 2 x 20 within subjects study was conducted where participants were asked to lie at random to five out of twenty questions. Resulting accuracy was slightly over 55% until a manipulation of baseline readings, which increased accuracy to 69%. This study had issues calculating reliable pupil diameter measurements due to the off-axis problem described by Andrew Duchowski et al. that results from participants eyes not gazing straight forward in relation to the HMD's eye tracking cameras.[1] This study provides a framework for lie detection in VR and reports promising findings that could be improved upon by further research. However, due to this study's failure to account for the off-axis problem, the results should be interpreted with the correct amount of skepticism.

KEYWORDS

Pupillary Dilation, Virtual Reality, VR, Lie Detection, Eye Tracking, Lie Detector, Eye Tracking.

1 INTRODUCTION

1.1 Goals

This experiment aims to use pupillary dilation as a means to indicate truth status. Since pupils are known to dilate under cognitive load, this can be used as a metric to guess if a person is lying.[2] This paper intends to study pupil dilation in a virtual reality setting and gather data to support that you can use pupil diameter as a metric for lie detection, specifically inside of VR.

Conference'17, July 2017, Washington, DC, USA

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

1.2 Motivations

Methods for determining whether a person is telling the truth or a lie have been proposed for centuries, as the ability to differentiate lies from deceit has tremendous legal implications. Techniques, such as the use of a polygraph machine, have been used extensively throughout modern history, however, some recent studies such as Saxe, et al.[3] have found the polygraph to be significantly less reliable than previously believed. While modern alternatives, such as fMRI-based lie detection show promising results, they are prohibitively expensive to use in most legal settings[4], and too complex in nature to explain to juries. Eye tracking could serve as a novel replacement for the polygraph in lie detection, and a more affordable option than fMRI. Furthermore, since pupil dilation is, by nature, heavily influenced by ambient lighting conditions, eye tracking in VR has the potential to eliminate this shortcoming by placing the subject in an environment with strictly controlled lighting. This paper aims to measure the effectiveness of such methods in detecting lies.

1.3 Hypothesis

Research has shown a correlation between a person's pupil dilation and their relative cognitive load. This has the potential to predict whether a subject's answer to a given question will be truthful or deceitful. Modern eye-tracking apparatuses can accurately and quickly measure the dilation of a person's pupil, however, they are subject to the changing lighting conditions of whatever room they are used in. VR-based eye tracking has the potential to outperform standalone eye-tracking apparatuses since lighting, background noise, and surroundings can be strictly controlled within the VR headset.

2 BACKGROUND

Lie detection technologies have been on a steady incline for over a century since the first device named the Lombrosso's Glove was invented in 1881.[5] Eye tracking capabilities net promising results in helping the validity of lie detection applications. The study of eye movements and pupil dilation can be monitored to help determine a user's truthfulness. Previous experiments have continued to validate the use of eye tracking in lie detection. A lie-detecting robot was created for a simulation of a human-robot interaction that studied the pupil diameter in a fixed environment to determine if the human was being truthful[6]. The ability to assess cognitive load to determine a lie stimulates unique opportunities for future

implementations. Another meta-analysis of lie detection used social cues to study the response to various questions answered by the participants to emphasize certain cues that had a bigger impact on the ability to successfully detect a lie. [7] The saccade (rapid eye movements), and pupil dilation showed accurate detection of lies. A pilot experiment resulted in significant differences when observing saccade amplitudes. During the truth-telling and lie-telling situations, it's observed that lying increases cognitive load. Therefore, cognitive load contributes to minimal eye movements and saccade amplitudes. [8]

3 METHODOLOGY

3.1 Apparatus

Figure 1: Researcher wearing HMD

This experiment used the Vive Pro Eye to capture eye-tracking data. Figure 1 shows a researcher wearing the HMD. The Vive Pro Eye has dual OLED 3.5" diagonal displays that operate at 1440 x 1600 pixels per eye. It has a 90 Hz refresh rate and a 110° Field of View. It comes with high-resolution audio as well. The Vive's eye trackers are binocular and have a 120 Hz refresh rate, an accuracy of 0.5-1.1°, and a 100° track-able field of view. The environment was created with the Unity game engine, which has a 90 Hz refresh rate. This experiment captured data through the Unity update system and opted for a 90 Hz refresh rate for collecting eye-tracking data.

3.2 Stimuli

This experiment places the participant in a virtual environment. This environment, presented in Figure 2, is a benign white room with a TV, nightstand, and potted plant. The participant will be prompted to focus on the TV. The TV will then guide them through the experiment.

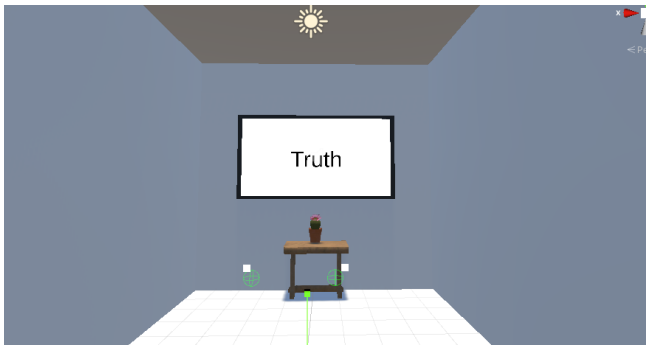

Figure 2: Virtual Environment

The participant will have audio played through the headphones of the Vive HMD. This audio will contain questions that the participant must answer and a guide telling them how the experiment will be run. There will be twenty questions and each question will be asked in less than five seconds. The TV in the room will play a variety of videos. It will play an introduction to the experiment where a researcher talks about what the experiment is and how it is run. It will also guide the participant through the study by prompting different actions in response to the twenty questions mentioned earlier. These actions are enumerated below.

(1) Figure 3 displays a prompt to tell the truth.

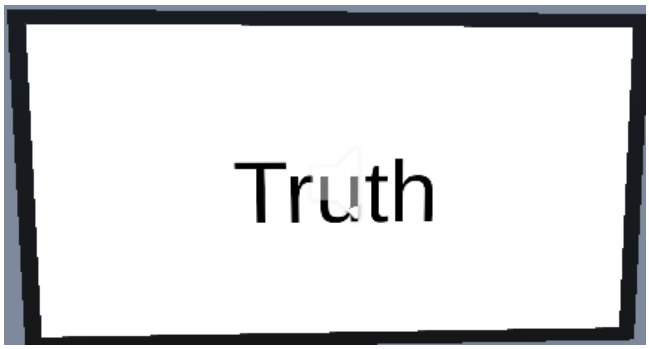

Figure 3: Prompt Participant to Tell Truth

(2) Figure 4 displays a prompt to tell a lie.
(3) Figure 5 displays a prompt indicating to the participant that a baseline reading is being collected.

3.3 Subjects

There are eleven subjects in this experiment. There were five male subjects ranging from age 18 to 35, and six female subjects ranging

Figure 4: Prompt Participant to Tell a Lie

Figure 5: Indicate to Participant Baseline Procedure

from age 18 to 54. Only one participant was over the age of 35. If they were excluded, it would make the female age range 18 to 28. These subjects were all local to the Clemson area. No subjects had any visual impairment.

3.4 Experimental Design

This experiment will use a 2x 20 within-subjects design. There are 20 different questions where the participant is asked to either lie or tell the truth. The lie and truth directions are chosen at random using selection sampling to ensure five questions are chosen as lies. Selected participants will listen to the same set of pre-recorded audio questions, and also be positioned in front of a virtual TV screen set in the virtual environment. Participants will be asked the same set of 20 questions (Table 1). Entering the questioning phase, the computer program will randomly assign instructions for each question to instruct the participant to either lie or tell the truth. Each question is designed to be answered with a simple true or false statement. Participants will audibly respond with either a lie or truth during the response phase. This methodology compares pupil dilation and the manner of prevarication because the audio stimuli were of the same set for all subjects should yield fair results. The independent variables of the experiment include the instructions given from the TV to lie or be truthful, and the questions asked. The same set of questions is asked to each participant, and all participants are in the same controlled virtual environment. The dependent variable would be the dilation difference during and after giving an audible response to the questions. The dependent variable could then be used to predict truthfulness.

Table 1: Questions asked to participants.

Question No.	Question
1	Two plus two equals five
2	Saturn orbits the moon
3	There are seventy five US states
4	Football is a contact sport
5	Orange is a mixture of red and yellow
6	A yard is longer than a meter
7	Clemson's colors are orange and grey
8	The English language has twenty eight letters
9	There are twelve inches in a foot
10	A lie is an opposite of a truth
11	Fall is the coldest season
12	Most vehicles have an automatic transmission
13	Rarity often has no effect on an item's value
14	Maserati is known for their watches
15	The human body cannot survive a week without water
16	Chocolate comes from a plant
17	Germany is not a part of Europe
18	Our oceans are our worlds most unexplored places
19	Cats are better than dogs
20	Soccer is the most popular sport on the planet

3.5 Procedures

Researchers will first introduce themselves to the participants, then give the participants a written description of instructions and the general aims of the experiment. Participants were given a short period to review the documentation and also fill out a survey provided to them by the researchers. Participants were then allowed to ask researchers any other questions they may have. The researchers then guided the participant through eye calibration on the Vive and started the study. When the study begins, participants are prompted via the TV in the room to inspect the virtual environment. After a short period, their attention is drawn back to the TV and a guide to the experiment is played on the TV. The guide walks through the study and informs the participant that the TV will prompt the user to lie or tell the truth. It then tells the participant that a short practice round will be played and then asks the participant if they are ready to start the real experiment. If so the real experiment begins, if not the same practice round is repeated. Once the real study has begun, the experiment will go through an array of questions. Questions will be chosen at random for the participant to lie about. Each question will consist of three phases. Phase one will be a baseline period. No audio will be played, and the TV will state that it is gathering baseline data. This period will last five seconds. Then phase two will begin. Phase two is the question-asking phase. During this phase, the participant will be prompted via the TV to lie or tell the truth while audio is played through the headphones asking a stimulus question. In phase three, participants will audibly answer yes or no to the provided question and either lie or tell the truth. This process will occur for each question and the study will end. At this point, the researchers removed the participant from the headset and asked if they had any further questions. Then an informal interview was conducted about the user's experience in the study.

4 DATA COLLECTION AND PREDICTION METHODOLOGY

4.1 Data Collection

At each refresh at 90Hz, the left and right pupil diameter was logged to a CSV file. When the eye was closed, -1 was reported as the pupil diameter. After the experiment was finished, data was processed by removing all -1 pupil diameter measures and interpolating from the value before the first measurement of -1 to the last measurement of -1 for each frame removed. In this way, pupil diameter is continuously available. For all predictions and figures, left pupil diameter was used. The resulting data appears akin to Figure 6. The data was then de-noised to give smoother data. This helps with the inconsistency of both pupil jitter and small inconsistencies with the eye tracker. Data was de-noised using the Savitsky Golay filter in Python's SciPy module. For the Savitsky Golay filter a window length of 41 samples was used with 1 for the polynomial order. When calculating the differential, the Savitsky Golay filter used a single derivation. De-noised data appears akin to Figure 7. Figure 8 shows a derivative of smoothed data. To help correct for eye movement throwing off the baseline calculations, the baseline mean and all other metrics were offset towards the proposed 3.25 mm pupil diameter average reported by Abdus Ansari et al. [9] The offset was the measured baseline mean and the 3.25 mm pupil metrics mean. After this offset, data would appear akin to Figure 9.

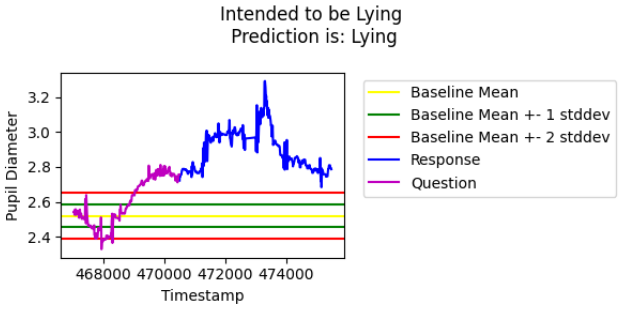

Figure 6: Un-smoothed Processed Data

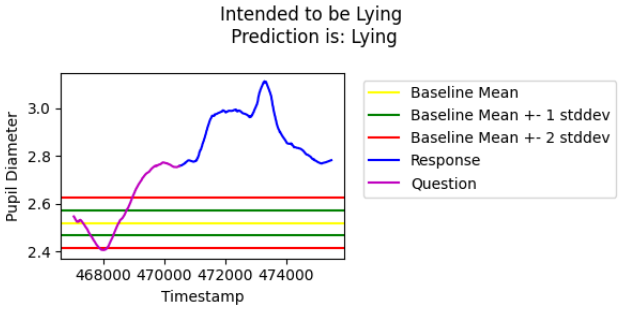

Figure 7: Smoothed Processed Data

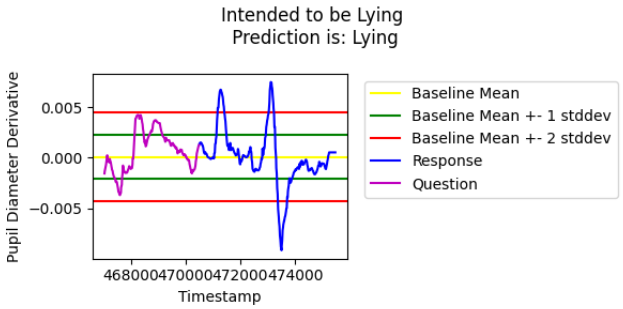

Figure 8: Smoothed Derivative

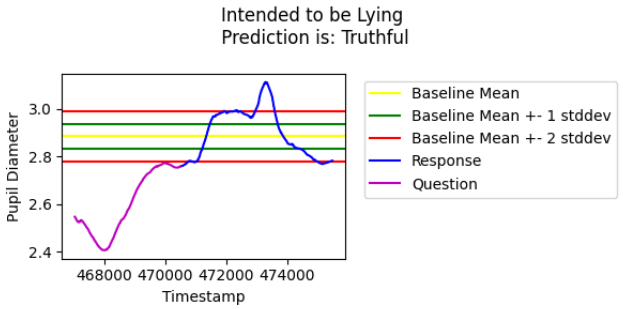

Figure 9: Smoothed Processed Data w/Offset

4.2 Making Predictions

Average pupil diameter was calculated during a baseline period before each question. This baseline mean was then used for that question and that question alone. The mean of the baseline portion and deviations up to two standard deviations are calculated. A person is predicted to be lying if, in the question and response portions of each question, the number of pupil diameter measurements taken that measure above two standard deviations of the baseline mean account for 25% of the question and response sections measurements.

5 RESULTS

Figure 10 contains a table full of results. Raw prediction results for each participant, with and without the 3.25 mm offset, are included in Tables 2 and 3, respectively.

5.1 Global Averages Over all Participants

This section describes the results from the left pupil diameter after being smoothed via the Savitsky Golay filter. The average total correct guesses per participant was 10.91 questions out of 20. Of the incorrect predictions, 8.45 of them were truths labeled as lies, and 0.64 of them was a lie labeled as a truth. The average prediction accuracy across all participants was 55%.

5.1.1 Best Participant Accuracy. The best predictions in an experiment contained 17 correct guesses with 1 false positive, but 2 false negatives. Prediction accuracy was 85%.

5.2 Global Averages Over all Participants w/Offset

In an attempt to reduce the number of false positives due to bad baseline calculations described in Section 6, an offset was applied to the participant data. Instead of using their baseline mean, we used their baseline mean averaged with the average pupil diameter in bright light conditions, 3.25 mm, described in the article by Abdus Ansari et al.[9] This resulted in slight improvements, resulting in an average prediction accuracy of 69%.

5.2.1 Best Participant Accuracy w/Offset. The best predictions in an experiment with an offset contained 19 correct guesses with 1 false positive, and 0 false negatives. Prediction accuracy was 95%.

6 PITFALLS WITH PROPOSED SOLUTIONS FOR FURTHER RESEARCH

Upon conducting exit interviews on subjects, some experimental oversights were made known.

6.1 Forced Fixation

6.1.1 The Problem of Looking Up. The intention of the HMD was that it would allow for less restrictive eye-tracking. In theory, since the HMD is attached to the user's head, we need not stabilize the head and make sure that it isn't moving in order to capture accurate pupil diameter metrics. However, in reality, an oversight was made. The participants often reported looking up when they needed to respond to a question or during the baseline calculation. This glance of the eye upward seems to substantially throw off the pupil diameter. As a result of a participant looking up, pupil diameter is reported to be smaller than in reality and false positives (predictions of a lie) are the result. When a participant looks straight forward, their eye appears as a circle, but when their eyes do not face forward, their eye appears as an ellipse. The diameter of the elliptic pupil shape is smaller and negatively impacts the prediction. This problem is dubbed the off-axis issue by Andrew Duchowski et al. [1]. A proposal for compensation is also made where subjects are directed not to move their eyes during the experiment.

6.1.2 The Problem of Body Movement. Similar to the last issue, some participants were very good at focusing on the intended fixation, the TV, but their bodies did not remain still. One subject constantly rocked left to right due to the chair having swivel functionality. As a result, their eyes with respect to the eye tracker did not stay centered. Due to this off-center pupil, the pupil appears smaller than in reality during baseline calculation and a false positive (lie prediction) is reported.

6.2 Poor Question Quality

6.2.1 Confusing Questions Due to Question Material. All but one participant hailed from the United States. Some questions involved the metric system and as a result, caused confusion among more than one participant. Since pupils respond to cognitive load, confusion will cause the pupil to dilate, and even if a participant is attempting to tell the truth the algorithm will predict a lie due to the mental strain of processing the question.
Another question involved the auto manufacturer Maserati and if they were known for making watches. Some participants reported not knowing what Maserati was or did. Removing questions that deal with brands would provide better results.

6.2.2 Confusing Questions Due to Question Style. Some questions were not purely factual. For instance, one question asked if cats were better than dogs. Some subjects were thrown off by the change of pace from asking a personal opinion and brought up this question specifically in the interview. No opinion response questions should be asked in future research without telling them to expect such questions.

7 SUMMARY AND CONCLUSION

In summary, this paper intended to investigate the validity of eye tracking in VR with the Vive Pro Eye HMD. Lie detection via pupil diameter measurement is plagued by factors that could potentially be mitigated in virtual environments. This study collected data from 11 participants who answered twenty questions of which five lies were chosen at random. Our method of using pupil diameter to predict lies resulted in 55% prediction accuracy across all participants. Researchers then further modified data to try to generate better results caused by experimental design mistakes mentioned in Section 6.1. by using established average pupil diameters to augment baseline readings.[9] This provided slightly better results, with a 69% accuracy over all participants. Further research is needed into lie detection via VR in order to make stronger claims on the methodology as well as iron out the experimental deficiencies of this paper.

REFERENCES

[1] A. T. Duchowski, K. Krejtz, I. Krejtz, C. Biele, A. Niedzielska, P. Kiefer, M. Raubal, and I. Giannopoulos, "The index of pupillary activity: Measuring cognitive load vis-à-vis task difficulty with pupil oscillation," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, (New York, NY, USA), p. 1–13, Association for Computing Machinery, 2018.

[2] K. Krejtz, A. T. Duchowski, A. Niedzielska, C. Biele, and I. Krejtz, "Eye tracking cognitive load using pupil diameter and microsaccades with fixed gaze," *PLOS ONE*, vol. 13, pp. 1–23, 09 2018.

[3] L. Saxe and G. Ben-Shakhar, "Admissibility of polygraph tests: The application of scientific standards post-daubert," *Psychology, Public Policy, and Law*, vol. 5, pp. 203–223, 03 1999.

[4] H. Hyman, "Functional mri as a lie detector: Its potential as evidence for or against malingering," *Plaintiff Magazene*, 2008.

[5] M. Vicianova, "Historical techniques of lie detection," *Europe's Journal of Psychology*, vol. 11, pp. 522–534, Aug. 2015.

[6] D. Pasquali, A. M. Aroyo, J. Gonzalez-Billandon, F. Rea, G. Sandini, and A. Sciutti, "Your eyes never lie: A robot magician can tell if you are lying," 2020.

[7] M. Hartwig and C. F. Bond Jr., "Lie detection from multiple cues: A meta-analysis," *Applied Cognitive Psychology*, vol. 28, no. 5, pp. 661–676, 2014.

[8] K. K. Lim, M. Friedrich, J. Radun, and K. Jokinen, "Lying through the eyes: Detecting lies through eye movements," p. 51–56, 2013.

[9] A. S. Ansari, J. Vehof, C. J. Hammond, F. D. Bremner, and K. M. Williams, "Evidence that pupil size and reactivity are determined more by your parents than by your environment," *Frontiers in Neurology*, vol. 12, 2021.

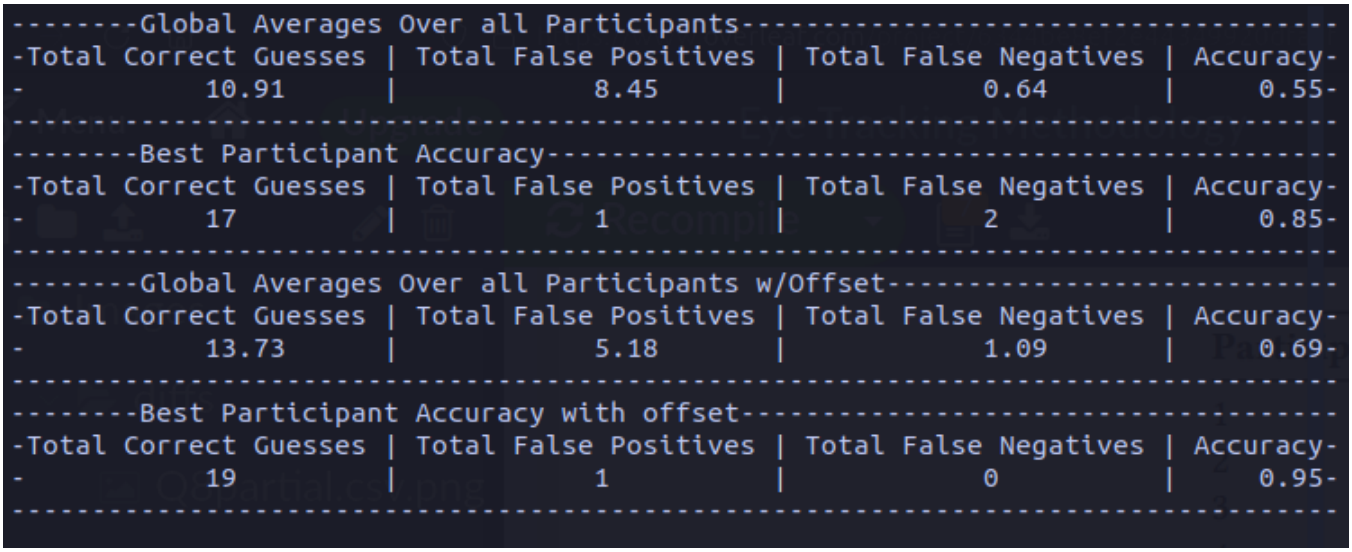

Figure 10: Final Results

Table 2: Raw Data, without offset

Participant	True Negative	False Positive	False Negative	True Positive
1	8	10	0	2
2	12	6	0	2
3	15	3	1	1
4	10	8	1	1
5	17	1	2	0
6	10	8	0	2
7	6	12	0	2
8	9	9	1	1
9	10	8	1	1
10	7	11	1	1
11	1	17	0	2
Total Accuracy				54.5%

Table 3: Raw Data, with offset

Participant	True Negative	False Positive	False Negative	True Positive
1	18	0	2	0
2	17	1	0	2
3	18	0	2	0
4	15	3	1	1
5	18	0	2	0
6	15	3	1	1
7	5	13	0	2
8	9	9	1	1
9	13	5	2	0
10	8	10	1	1
11	5	13	0	2
Total Accuracy				68.6%