Computer Vision and Image Understanding 147 (2016) 52–68

Contents lists available at ScienceDirect

Computer Vision and Image Understanding

journal homepage: www.elsevier.com/locate/cviu

Does "lie to me" lie to you? An evaluation of facial clues to high-stakes deception

Lin Su*, Martin Levine

Department of Electrical and Computer Engineering, Centre for Intelligent Machines, McGill University, 3480 University Street, MontrealQC H3A2A7, Canada

ARTICLE INFO

Article history:
Received 23 February 2015
Accepted 18 January 2016

Keywords:
High-stakes
Deception detection
Spontaneous facial expression
Unconstrained database
Random Forest

ABSTRACT

During a forensic interview, high-stakes deception is very prevalent notwithstanding the heavy consequences that may result. Studies have shown that most untrained people cannot perform well in discerning liars and truth-tellers.

Some psychological studies have stated that certain facial actions are more difficult to inhibit if the associated facial expressions are genuine. Similarly, these facial expressions are equally difficult to fake. This has cast light on the possibility that deception could be detected by analyzing these facial actions. However, to the best knowledge of the authors, there is no computer vision research that has attempted to discriminate high-stakes deception from truth using facial expressions. Therefore, this paper aims to test the validity of facial clues to deception detection in high-stakes situations using computer vision approaches.

We also note that only a limited number of the existing databases have been collected specifically for deception detection studies and none of them were obtained from real-world situations. In this paper we present a video database of actual high-stakes situations, which we have created using YouTube.

We have adopted 2D appearance-based methods as the methodology to characterize the 3D facial features. Instead of building a 3D head model as is the current trend, we have extracted invariant 2D features that are related to the 3D characteristic from nine separate facial regions by using dynamic facial analysis: eye blink, eyebrow motion, wrinkle occurrence and mouth motion. Then these cues are integrated to form a facial behavior pattern vector. A Random Forest was trained using the collected database and applied to classify the facial patterns into deceptive and truthful categories.

Despite the many uncontrolled factors (illumination, head pose and facial occlusion) contained in the videos in our database, we have achieved an accuracy of 76.92% when discriminating liars from truth-tellers using methods based on both micro-expressions and "normal" facial expressions. The results have shown that using facial clues for automated lie detection is very promising from the point of view of practice.

1. Introduction

Is it possible to step behind a person's facial "mask" to determine who is lying and who is telling the truth? Clearly, most people are incapable of doing this. Consider the recent TV show, called *Lie to Me*, which enjoyed a large popularity. In this TV series, Dr. Lightman and his team used their "talent" to assist the police with the investigation of criminal cases. Their "talent" was that they could visually determine whether a suspect was lying by interpreting his micro-expressions[1] during a face-to-face interrogation.

The deception of concern in *Lie to Me* is termed a *high-stakes deception*, because obviously, interrogation in a criminal case usually involves a high-stakes scenario. Different from lies told in our daily lives, deceptions in high-stakes situations are more likely to result in severe consequences. Therefore a liar in this situation might experience a heavy cognitive load, being aware of the penalties should his lie be detected. Considering the risk of releasing a guilty suspect or mistaking an innocent person, the detection of high-stakes deception is a necessity for a democratic society.

* Corresponding author.
E-mail address: lin.su3@mail.mcgill.ca (L. Su).

[1] A micro-expression is a rapid and involuntary facial expression, which can seemingly reveal a person's genuine emotion. It will be discussed in Section 3.

http://dx.doi.org/10.1016/j.cviu.2016.01.009
1077-3142/

Do humans have the ability to detect such high-stakes lies, just like Dr. Lightman and his colleagues did? The answer is no. Most untrained people are no better than chance at detecting lies [1], and the subjective decision-making process of humans will bias their decisions [2]. A study conducted in [3] showed that they were more likely to misidentify innocent people as liars.

However, regardless of the fact that ordinary people are deficient at deceit detection, we might ask whether the fundamental theory of detecting lies in *Lie to Me* is plausible. In other words, is a facial expression or micro-expression reliable as a clue to deception? This paper investigates this question to attempt to validate whether or not facial clues could be adopted as indicators of deception in high-stakes situations.

The proposed automated method consists of three stages of computer video analysis: pre-processing, dynamic feature analysis and classification. In the pre-processing stage, face detection and facial landmark localization are firstly applied to register the face. Then an anthropometric model is used to decompose the face into several facial regions. In the dynamic feature analysis stage, the indicators of deception, which are actually facial expressions, are detected in each facial region and collected into a facial behavior description vector. Finally in the classification stage, a binary Random Forest classifier is trained to discriminate deception and honesty.

The paper makes the following contributions:

1. It sets a precedent for future research on high-stakes deception detection using facial clues. To the best knowledge of the authors, no automated (or manual) computer vision methodology has shown the validity of facial expressions as indicators of high-stakes deception. As will be seen later in this paper, the results are very promising.
2. A database consisting of high-stakes deception videos of real-world situations has been collected from YouTube. In our database, the suspects are either pleading for the safe return of their missing relatives or denying their involvement in the disappearance or death of the victims. All of the criminal cases are real, and approximately half of the suspects were convicted as guilty by overwhelming evidence. As will be seen in Section 3.2, only a limited amount of the existing databases has been collected for deception detection studies, and none of them were obtained from *real-world* situations.
3. The proposed method is concerned with analyzing facial expressions in *unconstrained* environments. Since the videos in our database were collected from YouTube, certain uncontrollable factors added to the difficulty of their analysis. These totally unconstrained and spontaneous videos are subject to temporal variations in illumination, head pose and facial occlusion. However, in the current literature, little research has been conducted to address these issues with regard to facial expression analysis. Instead of seeking a solution for dealing with them, almost all of the studies to date have excluded certain data that were not ideal for the purpose of analysis. In comparison, the proposed method seeks to address the deception detection problem, but in the presence of exactly these factors.

The rest of this paper is organized as follows: Section 2 is a literature review of the background of the presented research, including an introduction to high-stakes deception, people's performance at discriminating liars and truth-tellers, current measures and automated methods for detecting lies. Section 3 presents the theoretical foundation of our deception detection method, the database collected by the authors, and proposes the dynamic feature analysis methods based on these theories. After feature extraction, the experimental procedures for training and testing a binary classifier to discriminate deception and honesty are presented. The results are discussed in Section 4. Finally the conclusion and future work are presented in Section 5.

2. Related work

2.1. People's performance at lie detection

People lie twice a day on average [4,5] but there are certain circumstances where lies are more likely to give rise to heavy consequences to both an individual and society [6]. These are termed *high-stakes situations*. Examples are being interrogated by a police officer, defending oneself in a courtroom, or appealing to a judge for a parole. However, interrogators rarely have the capacity to see the truth behind lies, regardless of the fact that they occur on a regular basis. Table 1 is a summary of the reported performance on deceit detection by *human observers*. In most cases, the lie detection accuracy of ordinary people has been shown to be only slightly above chance. Humans often believe in emotion, rather than empirically-based clues to deception, such as gaze aversion and fidgeting, which is detrimental to their judgment [7]. Interestingly emotionally intelligent people perform worse at deception detection, because they have greater sympathetic feelings for others [8]. Also, another factor is that people are often more skeptical of exonerating than incriminating witnesses [9], consequently degrading their performance at lie detection.

It is interesting that psychological studies have shown that a person's performance at lie detection can be improved by professional training. This has been attributed mainly to the fact that during the learning process, the measure of lying gradually changes from a naïve stereotype to appropriate empirical evidence.

2.2. Current deception detection measures

Evidence-based measures, including physiological, scientific, and psychological linguistic measures, have been applied to real world situations to discriminate liars and truth-tellers.

A physiological measure is a very straightforward way of distinguishing between guilty and innocent suspects in criminal cases. For instance, thermal body imaging [14,15] detects blood flow via special cameras, voice stress analysis [16] measures the voice signature, and the polygraph [17,18], which measures multiple physiological signals simultaneously. The drawback of these physiological approaches is that they only focus on measures emanating from the peripheral rather than central nervous system [18]. In this case, the interpretation of the measurements is not *directly* related to the emotional states of humans, which are controlled by the central nervous system. The other disadvantage is that they are easily degraded by certain countermeasures, such as tranquillizers or repeated practice.

As an alternative method of lie detection, functional Magnetic Resonance Imaging (fMRI) is actually an objective measurement of the mental state [19]. Typically, it employs an MRI scan of the head or body. On the theoretical side, fMRI is superior to the polygraph since it measures brain activity in the central rather than peripheral nervous system [20]. However, fMRI also has several shortcomings [19,21,22]. First, most of the fMRI studies have been small and rarely replicated, thereby lacking validity through scientific scrutiny. Second, brain activity varies considerably for each individual, making it unreliable to simply use the average brain activity of a group of healthy and normal people to evaluate the truthfulness of individual high-stakes cases. Third, the theoretical findings of many studies are not consistent with each other, due to differences in experimental settings and paradigms.

A psychological linguistic or verbal measure is another deceit detection approach which has been widely studied by