DREU Summer 2013 Final Report

Title: Deceptive Speech Across Cultures

Author: Nishmar Cestero

Abstract:

We define deception as the deliberate choice to mislead. We are collecting data for a large study of the features that characterize deceptive speech in different cultures and how people recognize deception. We hypothesize that acoustic, prosodic, and lexical cues can be used to identify deception automatically in American English, Arabic and Mandarin speakers (speaking English) more accurately than human judges. We propose that simple personality tests can predict individual differences in production and perception of deception across cultures, as we have found for American English speakers judging deception. Our goals are to develop automatic deception detection technologies and identify techniques to help select humans who are good at deception detection.

Background:

Deception is the deliberate choice to mislead; it is done without prior notification with the purpose of gaining some advantage or to avoid some penalty. Our definition excludes falsehoods due to ignorance or error, self-deception, delusions, or pathological behaviors, and theater.

The polygraph, which is used as a lie detector test, takes into account biometric factors, such as increases in blood pressure and perspiration. However it has been widely acknowledged that these factors are inadequate for deception detection. Previous deception research has focused on facial expression cues (Ekman '76, Frank '03), body posture and gestures (Burgoon et al '94), and brain imaging technologies (e.g. MRI) (Langleben et al '02). These methods are not optimal for general use because they require expensive and bulky equipment. Text-based approaches, which focus on variation in lexico-syntactic features (Streeter et al' 77) by suspects in their statements to authorities, are controversial and difficult to automate.

Previous efforts to study deceptive speech (Hirschberg et al '05) led to the creation of the CSC Corpus, the largest cleanly recorded audio corpus of within-subject deceptive speech. Automatic deception detection procedures performed up to 20% better than human judges. This method takes into account general lexical and acoustic/prosodic features used as cues to deception, as well as, subject dependent features, such as propensity to laugh. Additionally, interesting individual differences in some behaviors were observed. It was also found that human judges' accuracy in judging deception could be predicted from their scores on simple personality tests.

Goals:

DREU Summer 2013 Final Report

Our goals are to develop technologies to help humans detect deception by providing more relevant information, and to identify techniques to help select humans who are good at deception detection.

Hypotheses:

We hypothesize that acoustic, prosodic, and lexical cues can be used to identify deception in native Arabic and Mandarin speakers (speaking English) with accuracies greater than human judges. We propose that individual differences in these features could be predicted through simple personality tests; some personality traits will predict individual differences in production and perception of deceptive behavior across native and non-native English speakers, while some differences will be culture-specific or mediated by gender and culture of the deceiver and target.

Methodology:

We schedule participants in pairs for the experiment at the Speech Lab in Columbia University. We ask for their background information and administer the NEO-FFI personality test. We record a baseline test, where the experimenter asks casual questions about their everyday life. The three to four minutes of speech collected in the baseline will later be compared to the samples from the lying game, which will help determine subject dependent features. The main component of the experiment is a "lying game" that involves using a biographical questionnaire in the "fake resume" paradigm. It contains questions that range from casual to personal (e.g. 'Who ended your last romantic relationship?' or 'Have you ever owned a cat?'). During the game, participants have no visual contact and take turns interviewing each other using the biographical questionnaire. They are motivated by our payment scheme, which offers a chance to win additional money if they can successfully deceive their partner and make accurate judgments.

Current status and Future Work:

We are currently in the data collection phase of the experiment. Over 100 pairs of subjects have been recorded. This summer alone, we have recorded about 21 hours of speech from 70 subjects. We have extracted intensity measures for the sound files to check their quality. Our current participant pool includes American English and Mandarin Chinese speakers recruited from Columbia University and Barnard College campuses. Additionally, we are starting to correlate behavioral variation in lies vs. truth with standard personality test scores for speakers (NEO FFI).

We will soon be including Arabic speakers in our participant pool. Once we have enough data, we will begin feature extraction: Acoustic/Prosodic (i.e. duration, speaking rate, pitch, pause) and Lexico/Syntactic (i.e. laughter, disfluencies, hedges). We can then conduct machine learning experiments to identify features significantly associated with deceptive vs. non-deceptive speech.