Automatic Detection of Verbal Deception

Synthesis Lectures on Human Language Technologies

Editor
Graeme Hirst, ***University of Toronto***

Synthesis Lectures on Human Language Technologies is edited by Graeme Hirst of the University of Toronto. The series consists of 50- to 150-page monographs on topics relating to natural language processing, computational linguistics, information retrieval, and spoken language understanding. Emphasis is on important new techniques, on new applications, and on topics that combine two or more HLT subfields.

Automatic Detection of Verbal Deception
Eileen Fitzpatrick, Joan Bachenko, and Tommaso Fornaciari
2015

Semantic Similarity from Natural Language and Ontology Analysis
Sébastien Harispe, Sylvie Ranwez, Stefan Janaqi, and Jacky Montmain
2015

Learning to Rank for Information Retrieval and Natural Language Processing, Second Edition
Hang Li
2014

Ontology-Based Interpretation of Natural Language
Philipp Cimiano, Christina Unger, and John McCrae
2014

Automated Grammatical Error Detection for Language Learners, Second Edition
Claudia Leacock, Martin Chodorow, Michael Gamon, and Joel Tetreault
2014

Web Corpus Construction
Roland Schäfer and Felix Bildhauer
2013

Recognizing Textual Entailment: Models and Applications
Ido Dagan, Dan Roth, Mark Sammons, and Fabio Massimo Zanzotto
2013

Linguistic Fundamentals for Natural Language Processing: 100 Essentials from Morphology and Syntax
Emily M. Bender
2013

Semi-Supervised Learning and Domain Adaptation in Natural Language Processing
Anders Søgaard
2013

Semantic Relations Between Nominals
Vivi Nastase, Preslav Nakov, Diarmuid Ó Séaghdha, and Stan Szpakowicz
2013

Computational Modeling of Narrative
Inderjeet Mani
2012

Natural Language Processing for Historical Texts
Michael Piotrowski
2012

Sentiment Analysis and Opinion Mining
Bing Liu
2012

Discourse Processing
Manfred Stede
2011

Bitext Alignment
Jörg Tiedemann
2011

Linguistic Structure Prediction
Noah A. Smith
2011

Learning to Rank for Information Retrieval and Natural Language Processing
Hang Li
2011

Computational Modeling of Human Language Acquisition
Afra Alishahi
2010

Introduction to Arabic Natural Language Processing
Nizar Y. Habash
2010

Cross-Language Information Retrieval
Jian-Yun Nie
2010

Automated Grammatical Error Detection for Language Learners
Claudia Leacock, Martin Chodorow, Michael Gamon, and Joel Tetreault
2010

Data-Intensive Text Processing with MapReduce
Jimmy Lin and Chris Dyer
2010

Semantic Role Labeling
Martha Palmer, Daniel Gildea, and Nianwen Xue
2010

Spoken Dialogue Systems
Kristiina Jokinen and Michael McTear
2009

Introduction to Chinese Natural Language Processing
Kam-Fai Wong, Wenjie Li, Ruifeng Xu, and Zheng-sheng Zhang
2009

Introduction to Linguistic Annotation and Text Analytics
Graham Wilcock
2009

Dependency Parsing
Sandra Kübler, Ryan McDonald, and Joakim Nivre
2009

Statistical Language Models for Information Retrieval
ChengXiang Zhai
2008

© Springer Nature Switzerland AG 2022
Reprint of original edition © Morgan & Claypool 2015

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopy, recording, or any other except for brief quotations in printed reviews, without the prior permission of the publisher.

Automatic Detection of Verbal Deception

Eileen Fitzpatrick, Joan Bachenko, and Tommaso Fornaciari

ISBN: 978-3-031-01030-9 paperback
ISBN: 978-3-031-02158-9 ebook

DOI 10.1007/978-3-031-02158-9

A Publication in the Springer series
SYNTHESIS LECTURES ON HUMAN LANGUAGE TECHNOLOGIES

Lecture #29
Series Editor: Graeme Hirst, *University of Toronto*
Series ISSN
Print 1947-4040 Electronic 1947-4059

Automatic Detection of Verbal Deception

Eileen Fitzpatrick
Montclair State University

Joan Bachenko
Linguistech LLC

Tommaso Fornaciari
Italian National Police

SYNTHESIS LECTURES ON HUMAN LANGUAGE TECHNOLOGIES #29

ABSTRACT

The attempt to spot deception through its correlates in human behavior has a long history. Until recently, these efforts have concentrated on identifying individual "cues" that might occur with deception. However, with the advent of computational means to analyze language and other human behavior, we now have the ability to determine whether there are consistent clusters of differences in behavior that might be associated with a false statement as opposed to a true one. While its focus is on verbal behavior, this book describes a range of behaviors—physiological, gestural as well as verbal—that have been proposed as indicators of deception. An overview of the primary psychological and cognitive theories that have been offered as explanations of deceptive behaviors gives context for the description of specific behaviors. The book also addresses the differences between data collected in a laboratory and "real-world" data with respect to the emotional and cognitive state of the liar. It discusses sources of real-world data and problematic issues in its collection and identifies the primary areas in which applied studies based on real-world data are critical, including police, security, border crossing, customs, and asylum interviews; congressional hearings; financial reporting; legal depositions; human resource evaluation; predatory communications that include Internet scams, identity theft, and fraud; and false product reviews. Having established the background, this book concentrates on computational analyses of deceptive verbal behavior that have enabled the field of deception studies to move from individual cues to overall differences in behavior. The computational work is organized around the features used for classification from *n*-gram through syntax to predicate-argument and rhetorical structure. The book concludes with a set of open questions that the computational work has generated.

KEYWORDS

credibility assessment, deception detection, factual language, forensic linguistics, gold-standard data, ground truth, high-stakes scenarios, imaginative language, real-world data, stylometry, text classification

Eileen Fitzpatrick: To my husband, Ralph Grishman

Joan Bachenko: In memory of my mother, Claire Joan Baumgartner

Tommaso Fornaciari: In memory of my father, Alberto Fornaciari

Contents

Preface xv

Acknowledgments xvii

1 Introduction 1

1.1 Introduction 1
1.2 Verbal Cues to Deception 2
1.2.1 Linguistic Features Used in Identifying Deception 4
1.2.2 Effectiveness of Linguistic Cues to Deception 5
1.2.3 Verbal Cues to Ground Truth 5
1.3 What's Ahead 6

2 The Background Literature on Behavioral Cues to Deception 7

2.1 Introduction 7
2.2 Nonverbal Cues to Deception 7
2.2.1 Polygraphy 9
2.2.2 Voice Analysis: VSA and LVA 10
2.2.3 Thermography 11
2.2.4 Brain Scan: EEG and MRI 11
2.2.5 Vocal Cues 12
2.2.6 Body and Facial Movements 13
2.3 The Psychology Literature 13
2.3.1 DePaulo et al.'s Study 13
2.3.2 Vrij's Studies 28
2.4 The Forensic Literature 32
2.4.1 Statement Analysis 32
2.4.2 Statement Validity Analysis 33
2.4.3 Reality Monitoring 35
2.5 Forensic Implementations of the Literature 36
2.5.1 SCAN as an Investigative Tool and Training Program 37
2.5.2 Evaluations of SCAN 37

3 Data Sources . . . 43

3.1 Introduction . . . 43
3.2 Establishing Ground Truth . . . 44
3.2.1 Forensic Data Sources: Spoken and Written . . . 45
3.2.2 Financial Reports . . . 46
3.2.3 Mass Media Communications . . . 47
3.3 Risks with Ground Truth Sources . . . 48
3.3.1 Legal and Forensic Interviews and Statements . . . 49
3.3.2 Financial Reports . . . 50
3.3.3 Mass Media Communications . . . 51

4 The Language of Deception: Computational Approaches . . . 53

4.1 Computational Approaches to Verbal Deception . . . 53
4.1.1 Establishing Comparative Measures of System Performance . . . 53
4.1.2 Classification and Ranking . . . 54
4.1.3 Training and Testing . . . 55
4.1.4 System Evaluation . . . 55
4.1.5 Prepping the Data . . . 56
4.2 Considerations Specific to Deception . . . 56
4.2.1 Data Types Amenable to Deception Research . . . 57
4.2.2 Unit of Analysis: The Liar or the Lie . . . 57
4.2.3 Lies of Omission and Commission . . . 58
4.2.4 Level of Data Used for Modeling . . . 59
4.2.5 Training Data and Ground Truth . . . 59
4.3 The Current Systems . . . 59
4.3.1 Characters and *n*-grams . . . 60
4.3.2 Features . . . 66
4.3.3 Studies Looking above the Lexical Level . . . 75
4.4 Conclusion . . . 80

5 Open Questions . . . 81

5.1 Introduction . . . 81
5.2 Impact of Contextual Factors on Deceptive Narrative . . . 81
5.3 Deceptive Language and Imaginative Language . . . 82
5.4 Measuring the Distance between Diverse Narratives . . . 82
5.5 Ground Truth Annotation: The Search for Gold-Standard Data . . . 83

5.6 A Common Data Set 85
5.7 Cue Clustering 85
5.8 Correlation of Verbal with Nonverbal Cues 86
5.9 Conclusion 87

Bibliography 89

Authors' Biographies 101

Preface

There are many venues where the ability to spot the lie is an important, often critical, skill. It is necessary in police, security, border crossing, customs, and asylum interviews; in congressional hearings; in financial reporting; in legal depositions; in human resource evaluation; in predatory communications, including Internet scams, identity theft, and fraud; and in false advertising. Discovering lies can thwart serious immediate threats, provide productive directions in the investigation of past events, and assist in the accurate prediction of likely future events.

For much of the 20th century, the fields of psychology and criminal justice have studied the behaviors that might be associated, directly or indirectly, with deception. Three types of behavior have been examined: physiological behavior; vocal behavior, including prosodic features; and verbal behavior, including the words and structures that might correlate with deception.

The study of verbal behavior in deception is relatively new and the attention that natural language processing has paid to discriminating true from false claims is even newer, with most of the work done in the last 10 years as classification techniques have improved. Now is a good time to review the prior literature on deception and consider the NLP approaches that have been tried. Knowing the foundations and trends in work on deception, both theoretical and applied, will enable us to move forward productively.

Several areas of NLP are ripe to address the vocal and verbal features that might be associated with deception and new approaches may well combine information from these with the facial and body movements associated with deception. A spate of recent NLP papers on the classification of narratives as truthful or deceptive suggests that the field is ready to open up to this promising area.

The genesis of this text was a workshop on deception detection that took place as part of the European Meeting of the Association for Computational Linguistics in Avignon in the spring of 2012. The workshop brought together 35 colleagues and offered 14 presentations on work that ranged from annotation tools and corpus building for deception to cross-linguistic classification of deceptive narrative. It also gave the authors of this text an opportunity to work together on our common interest in the use of "real-world" data, primarily legal data, in NLP deception studies.

Following the mission of the Synthesis series, the present text is designed to give the student or researcher in natural language processing a background in the history of deception studies concentrating on the behavioral cues to deception that have been supported in the psychology, applied psychology, and criminal justice literature, a consideration of the real-world data sources for NLP work in deception, a review of NLP work in deception organized around the features

used for classification from ngram to rhetorical structure, and a look at exciting questions and areas that need to be addressed in order for the field to progress.

Eileen Fitzpatrick, Joan Bachenko, and Tommaso Fornaciari
July 2015

Acknowledgments

We wish to thank the students who took part in a graduate seminar on current issues in NLP at Montclair State University in which a previous draft of this text was used as the primary text for the course, including Richard Barrett, Janette Martinez, Matthew Mulholland, and Emily Olshefski. Their ability to take the material covered in the text to the next step in their own research gives us confidence that we are respecting the mission of this series.

We would also like to thank Ralph Grishman and Massimo Poesio for input, checking, and fruitful discussions related to the issues in this book.

We are indebted to Myle Ott and another, anonymous, reviewer who provided detailed, thoughtful feedback in their reviews of a draft of the book, and to the series editor Graeme Hirst who first suggested a text on deception in Avignon and has guided it through, with great patience and care, to reality. We are also grateful to the many authors whose work we have cited here.

Eileen Fitzpatrick, Joan Bachenko, and Tommaso Fornaciari
July 2015