Detecting Concealed Information in Text and Speech

Shengli Hu, Cornell University

Motivation

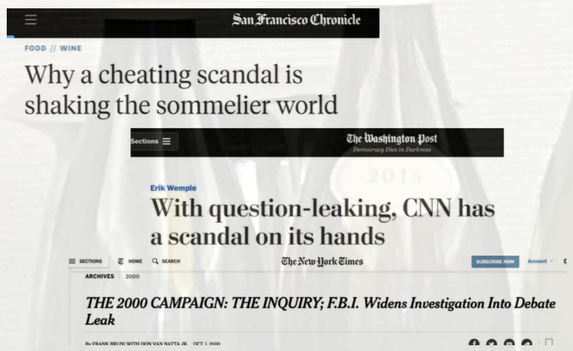

Research Questions

- How good are humans at detecting concealed information in technical settings?
- Can we improve on human performance?
- How are indicators of concealed information related to those of deception?
- When are Machine Learning classifiers better(or worse) than human domain experts?

Contributions

- The first corpus and study on concealed information in technical settings – please let me know if I am totally wrong here!
- Novel insights from identified key features (cf. deception)
- Multi-task learning framework with acoustic-linguistic features

Summary

- Acoustic-prosodic indicators appear largely consistent with deception
- Linguistic cues appear largely the opposite of deception
- Algorithms outperform domain experts by over 15%
- Multi-task learning framework with acoustic and linguistic features outperform baseline by over 11%

Selected References

Wolfgang Ambach, et al. 2010. A concealed information test with multimodal measurement.
Guozhen An, et al. 2018. Deep personality recognition for deception detection.
Stefan Benus, et al. 2006. Pauses in deceptive speech.
Sarah Ita Levitan et al. 2016.Combining acoustic-prosodic, lexical, and phonotactic features for automatic deception detection.
Sarah Ita Levitan, et al. 2015. Individual differences in deception and deception detection.
Sarah Ita Levitan, et al. 2018a. Acoustic-prosodic indicators of deception and trust in interview dialogues.
Sarah Ita Levitan, et al. 2018b. Linguistic cues to deception and perceived deception in interview dialogues.
Gideon Mendels, et al. 2017. Hybrid acoustic-lexical deep learning approach for deception detection.
Rada Mihalcea and Carlo Strapparava. 2009. The liedetector: Explorations in the automatic recognition of deceptive language.
Esther Mobley. 2018. Why a cheating scandal is shaking the sommelier world.
Myle Ott, et al. 2011. Finding deceptive opinion spam by any stretch of the imagination.

Deception vs. Information Concealment

The Information Grid		Appearance	
		Information	No Information
Truth	Information	Honesty	Concealed Information
	No Information	Deception	Honesty

- Audio recordings of 49 blind tasting practice sessions with 41 (in total) certified or advanced sommeliers;
- Written answer sheets of descriptors, calls and guesses;
- Demographics: gender, native language, wine credential, self confidence.

- Disclaimer: the author herself is a WSET diploma student, certified specialist of wine, certified sommelier, and certified specialist of spirits.

Outline

Corpus Collection

⇓

Preprocessing, Annotation, Feature Extraction

⇓

Empirical Analysis of Feature Sets for Significant Differences

⇓

Individual Differences in Concealing Information

⇓

ML/DL Classification from Text and Speech

Model	Features	F1 (single / multiple turn)
Logistic Regression	Bigrams	Human: NA / 56.28 61.18 / 65.45
Random Forest	IS 2009	59.23 / 60.03
MLP	IS 2009	63.96 / 67.27
BiLSTM	GloVe	61.41 / 67.35
MLP + BiLSTM	IS 2009, GloVe	64.12 / 68.57
MLP + BiLSTM	IS 2009, Individual Features, GloVe.	64.14 / 70.02
MLP + BiLSTM + Multi-task	IS 2009, Individual Features, GloVe.	65.16 / 71.51

Consistent or Inconsistent with most recent deception literature[2]

Feature	Concealed Information	Truthful
N-grams	*yeah*, but it, citrus, correct, ruby, did not, lift, botrytis, would not	*uh um*, there is, there are, was like, so, slight, not sure, blossom, clear
LIWC	*clout*, certain, function, cogproc, negate, *discrep*, differ, assent, posemo	compare, pronoun, verb, *ingest*, feel
Syntax	*adj*, *adverb*, syn_distinct	
Else	*specificity*, Δ(Trans, Text)	hedging[3], #word, length

Statistical significant indicators of concealed information

Feature	Male	Female	Low Skill	High Skill	All
Pitch (max)	S				S
Pitch (mean)					
Intensity (max)	S	S	(S)		S
Intensity (mean)		(S)			
Speaking Rate			S		S
Duration		(-)(S)	(-)(S)		(-)(S)
Voice Quality