Expert Systems with Applications 38 (2011) 8392–8398

Contents lists available at ScienceDirect

Expert Systems with Applications

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

An investigation of data and text mining methods for real world deception detection

Christie M. Fuller [a], David P. Biros [b], Dursun Delen [c,*]

[a] College of Business, Louisiana Tech University, P. O. Box 10318, Ruston, LA 71272, United States
[b] Spears School of Business, 415 Business Bldg, Oklahoma State University, Stillwater OK 74078, United States
[c] Spears School of Business, Oklahoma State University, Department of Management Science and Information Systems, 700 North Greenwood Avenue, North Classroom Building #431, Tulsa, OK 74106, United States

ARTICLE INFO

Keywords:
Deception detection
Data mining
Text mining
Information fusion
Classification
Credibility assessment

ABSTRACT

Uncovering lies (or deception) is of critical importance to many including law enforcement and security personnel. Though these people may try to use many different tactics to discover deception, previous research tells us that this cannot be accomplished successfully without aid. This manuscript reports on the promising results of a research study where data and text mining methods along with a sample of real-world data from a high-stakes situation is used to detect deception. At the end, the information fusion based classification models produced better than 74% classification accuracy on the holdout sample using a 10-fold cross validation methodology. Nonetheless, artificial neural networks and decision trees produced accuracy rates of 73.46% and 71.60% respectively. However, due to the high stakes associated with these types of decisions, the extra effort of combining the models to achieve higher accuracy is well warranted.

1. Introduction

The volume of text-based chat, instant messaging, and text-messaging, as well as the number of text-based on-line communities of practice are rapidly increasing. Even email continues to grow in use. For many companies today, email is a critical asset and the loss of that tool would bring them to a standstill. With the massive growth of text-based communication, the potential for people to deceive through computer-mediated communication has also grown and such deception can have disastrous results. Consider the example of the young California girl who was deceived into believing that her boyfriend thought the world would be better off with her dead, leading her to commit suicide. It turned out that the senders of the messages were another teenage girl and that girl's mother (Indictment puts internet pranksters on notice, 2008). People have been duped into believing fraudulent financial schemes, children have been coaxed into meetings with predators, and hacking exploits like "phishing" have compromised security. Such cases underscore the vulnerabilities related to text-based deception and information manipulation.

Unfortunately, humans tend to perform poorly at deception detection tasks in general. This phenomenon is exacerbated in text-based communications. A large part of the research in the detection of deception (also known as credibility assessment) has involved face-to-face meetings and interviews. Yet, with the growth of text-based communication, text-based deception detection techniques are essential.

2. Background

Techniques for successfully detecting deception, or lies, have wide applicability. Law enforcement can use these decision support tools and techniques to investigate crimes, conduct security screening in airports, or monitor communications of suspected terrorists. Human resources professionals might use deception detection tools for applicant screening. These tools also have the potential to screen emails to uncover fraud or other wrongdoings committed by corporate officers. While some may believe they can readily identify those who are not being truthful, a summary of deception research showed that on average, people are only 54% accurate in making veracity determinations (Bond & DePaulo, 2006). For the purposes of this study, deception is defined as "a message knowingly transmitted by a sender to foster a false belief or conclusion by the receiver" (Buller & Burgoon, 1996). The 54% figure may actually worsen when humans try to detect deception in text. In one study, people could only find about one-third of the lies that had been planted in text (George & Keane, 2006). The sheer volume of text-based communication can only further complicate the task of deception detection for humans.

The potential uses for a text-based deception detection tool are widespread, as are the potential sources of text to be analyzed. It is not surprising that a previous study has shown that accompanying an increased use of text communication, such as email and instant

* Corresponding author. Tel.: +1 918 594 8283; fax: +1 918 594 8281.
E-mail addresses: cfuller@latech.edu (C.M. Fuller), david.biros@okstate.edu (D.P. Biros), Dursun.Delen@okstate.edu (D. Delen).

0957-4174/$ - see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.eswa.2011.01.032

messaging, is a tendency to lie in these forms of communication (Hancock, Thom-Santelli, & Ritchie, 2004). In addition to text that originates in a typed form such as emails, instant messages, or word processing documents, text-based deception detection can be applied to any verbal communication that can be reproduced in typed, electronic form. For example, handwritten documents can be transcribed, as was done in this study. Voice recognition software or other means of transcription can be used to record oral communications, which could then be evaluated.

In addition to its wide applicability, text-based deception detection also overcomes some of the disadvantages of the well-known polygraph tool. While the polygraph is a fairly accurate tool for evaluating non-verbal indicators of deception, it has several drawbacks. It requires the presence of both a trained examiner and polygraph equipment. A structured interview must be conducted, during which the equipment must be attached to the interviewee in five ways (Twitchell, Jensen, Burgoon, & Nunamaker, 2004; Vrij, 2000). Persons being interviewed are always aware that their veracity is being judged, thus making it possible for them to force an inconclusive test. Also, in some parts of the world, cultural and religious norms may make it difficult to conduct a polygraph examination. For example, in some Muslim communities it may be culturally inappropriate for a male polygraph analyst to connect the leads of the device to a female subject. Further, certain questions may also be deemed to be inappropriate. In such conditions, the polygraph may be a sub-optimal solution.

Recently, a handheld version of the polygraph, known as the Preliminary Credibility Assessment Screening System (PCASS) was introduced in the field in Afghanistan by the US Military (Battelle Memorial Institute, 2007; Harris & McQuarrie, 2009; Senter, Waller, & Krapohl, 2006). The portable PCASS is easy to use and provides quick results, overcoming some of the shortcomings of the conventional polygraph tool. This is certainly an exciting development in the field of credibility assessment. However, using this device still requires that it be attached to the person of interest during a structured interview and it suffers from many of the same obstacles as the traditional polygraph.

The PCASS has not been subjected to rigorous testing in the field and appears to be less accurate than the polygraph, which has been reported to be 72–91% accurate in field studies (National Research Council, 2003). One PCASS study reports an overall accuracy of nearly 79%, though in that case the tool was trained with a sample containing over 80% deceptive cases (Harris & McQuarrie, 2009). The tool was 86% accurate in correctly identifying deceptive cases, but only 50% accurate in identifying truthful cases. The PCASS is about 62–63% accurate when the tool was trained and tested on samples containing about equal proportions of truthful and deceptive statements, termed a 'balanced sample' (Battelle Memorial Institute, 2007; Senter et al., 2006). While tools like the traditional polygraph and PCASS can aid investigators in assessing credibility based on nonverbal responses, they were never designed to evaluate other forms of communication, especially text-based communication.

The text-based deception detection tool studied here offers many of the same advantages as the handheld polygraph tool, namely portability, ease of use, and fast results. The text-based tool offers the additional advantages that it does not need to be physically attached to the person of interest and it does not require a structured interview. It only requires that a text-sample be captured in some manner, and can evaluate many text samples in a very short period. Results to date show that text-based deception detection is also more accurate than the PCASS in laboratory studies when trained on a balanced sample (Twitchell et al., 2006; Zhou, Burgoon, Twitchell, Qin, & Nunamaker, 2004).

Given the importance of the situations in which these tools might be used, accuracy is quite important. The tools must be shown to work accurately in samples that resemble the actual circumstances in which the tools will be used. Most research in deception detection in general and also specifically in text-based deception detection has used samples from student subjects in laboratory settings (DePaulo et al., 2003; Vrij, 2000), though a need for research using serious, or high-stakes, lies has also been identified (DePaulo et al., 2003; Frank & Feeley, 2003; Kohnken, 1985). While these results have established a solid research foundation, it is difficult to replicate the severity of real-world situations such as criminal investigations, in the laboratory. Polygraph research has illustrated this, showing that results differ for polygraphs of those involved in 'mock crimes' and field studies of those involved in actual crimes (Pollina, Dollins, Senter, Krapohl, & Ryan, 2004). Samples generated from mock crimes also suffer from a lack of consequence. The mock deceiver has nothing to fear if detected. Such differences in findings are likely to exist using other deception detection techniques as well. Since the stakes in real world deception detection can be quite high (i.e. incarceration, etc), error is unacceptable. Therefore, testing deception detection techniques in real world conditions is imperative. The sample analyzed here is text produced by those involved in actual crimes on military bases, allowing us to study deception detection in a real-world, high-stakes environment where serious consequences were possible.

3. Methodology

This study analyzed person-of-interest statements completed by people involved in crimes on military bases. In these statements, suspects and witnesses are required to write their recollection of the event in their own words. Base law enforcement (LE) personnel searched archival data for statements that they could conclusively identify as being truthful or deceptive. These decisions were made on the basis of corroborating evidence and case resolution (i.e. not just the personal opinion of LE personnel). The definition of deception relies on an intentional communication of false information, therefore statements where a person-of-interest was simply mistaken in their recall of events were not labeled as deceptive. Once labeled as truthful or deceptive, the law enforcement personnel removed identifying information and gave the statements to the research team. In total, 371 statements were used in our analysis. The statements were from many different types of crimes, such as traffic infractions, shoplifting, assault, and arson. All statements were provided by adults.

We underscore the importance of this data set. Unlike many past studies that used data collected from student groups conducting mock lies or deceptions, the individuals involved in these cases faced severe consequences for lying on an incident statement. Military members could face penalties up to and including courts martial for creating a false official statement. Civilians could face disbarment from the base, or in the case of DoD employees, termination of employment. These penalties are of course in addition to those that the person may be facing due to conviction for involvement in the crime.

3.1. Message feature mining process

This automated text-based deception detection method is based on a process known as message feature mining (MFM) (Adkins, Twitchell, Burgoon, & Nunamaker, 2004). This process relies on elements of data mining and text mining techniques. Traditionally, data mining analyzes categorical or numerical variables to find meaningful patterns in a large volume of structured/tabular data (Berry & Linoff, 2004). Text mining also seeks to find meaningful patterns in data, though the data usually originates as unstructured text. This text must be transformed into some structured format prior to analysis (Feldman & Sanger, 2007). As shown in Fig. 1, both data mining and text mining techniques are incorporated into the process. The overall process begins with preparing the data for processing. Here, the statements were originally handwritten and each had to be transcribed into a word processing file.

Next, the features, or cues, were determined. Over 30 different linguistic features have been previously identified (Bond & Lee, 2005; Hancock, Curry, Goorha, & Woodworth, 2005; Zhou, Burgoon, Twitchell, 2004) that may help differentiate between truthful and deceptive speakers. These cues originate from the theories used to study deception, including: Reality Monitoring, Interpersonal Deception Theory, Information Manipulation Theory, and the Self-Presentational Perspective of deception. Reality Monitoring was not originally developed as a theory of deception, but has been extended to this context. Reality Monitoring theorizes that memories based on actual experiences versus memories based on imagined events are distinct on several dimensions (Johnson & Raye, 1981). The Self-Presentational Perspective of deception proposes five ways in which deception may be revealed: liars are less forthcoming than truthtellers, liars will tell less compelling tales, liars will be less positive and pleasant, liars will be more tense and liars will include fewer ordinary imperfections and unusual contents within their messages (DePaulo et al., 2003). Information Manipulation Theory (IMT) proposes that deceptive messages violate conversational maxims of quality, quantity, relation and manner (McCornack, 1992). IDT views deception as an interactive form of communication, merging the principles of deception with those of interpersonal communication (Buller & Burgoon, 1996). Though originally developed for the study of deception in richer media, such as face-to-face communication, later work has suggested that IDT is applicable for studying most forms of communication. (Zhou, Burgoon, Nunamaker, Jay, & Twitchell, 2004). Based on these theories, it is believed there will be variation in the level of cues expressed between truthful and deceptive speakers.

For example, truthful speakers may include more details related to space, senses and time since these details are available to someone describing a real experience. Deceivers, on the other hand, may speak more about cognitive processes as they are having to imagine and artificially create the experience. The features included in this study represent categories or types of language indicators that are relatively independent of the text content and that can be readily analyzed by automated means. For example, first-person pronouns such as 'I' or 'me' can be identified without analysis of the surrounding text. A list of all features used in this study (along with their short descriptions) is shown in Table 1.

The statements were processed to determine the presence of the various features using a combination of two software packages: General Architecture for Text Engineering (GATE) and Linguistic Inquiry and Word Count (LIWC) (Cunningham, 2002; Pennebaker & Francis, 2001). These tools use dictionaries to identify and calculate the desired features.

After the features were identified within the statement, the value of each variable/cue was determined. For example, a part-of-speech tagger identifies the verbs in each statement, and then counts those verbs. For word, sentence and verb quantity, the result is a count of each feature for each statement. For the rest of the cues, a ratio of the words belonging to a given category to total words in the statement was calculated. The results, in a flat file format, were the input for the next step in the process, the feature selection.

3.2. Feature selection

One of the most common tasks in this type of predictive data mining study is to select the most appropriate features (i.e., cues) from a long list of candidates. This is especially true when the size of the data sample (number of person-of-interest statements in this study) is relatively small. The feature selection algorithm used in this study computes a *Chi-square* statistic for each feature. For continuous features, the algorithm divides the range of values in each predictor into *k* intervals (10 intervals is commonly used as a default; to "fine-tune" the sensitivity of the algorithm to different

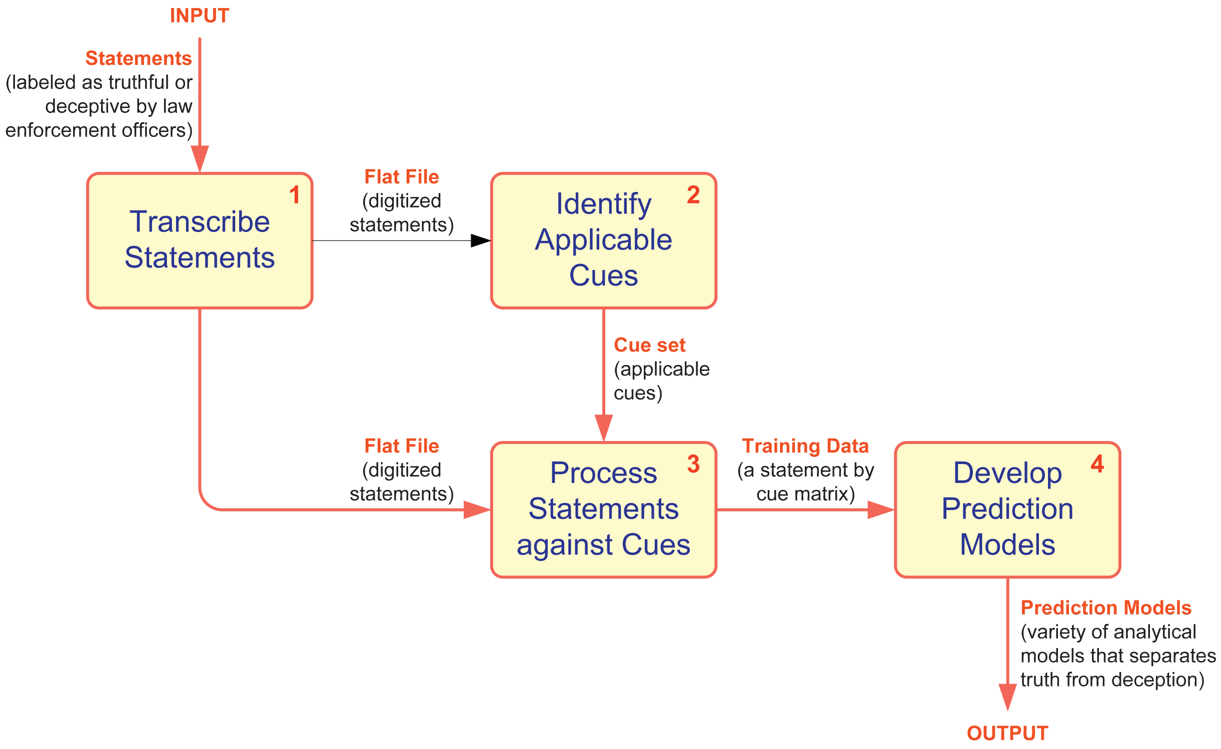
INPUT
Statements
(labeled as truthful or deceptive by law enforcement officers)
1
Transcribe Statements
Flat File
(digitized statements)
2
Identify Applicable Cues
Cue set
(applicable cues)
Flat File
(digitized statements)
3
Process Statements against Cues
Training Data
(a statement by cue matrix)
4
Develop Prediction Models
Prediction Models
(variety of analytical models that separates truth from deception)
OUTPUT

Fig. 1. Message feature mining process.

Table 1
Text-based deception features used in this study.

Input features (the independent variables):

1 1st person plural pronouns: Pronouns such as *we*, *our*.
2 1st person singular pronouns: Pronouns such as *I*, *me*, *mine*.
3 2nd person pronouns: Pronouns such as *you*, *your*.
4 3rd person pronouns: Pronouns such as *he*, *she*, *they*.
5 Activation: Extent to which language is active or passive.
6 Average sentence length: The average number of words per sentence.
7 Average word length: The average number of letters per word.
8 Bilogarithmic type-token ratio: This is a measure of unique words adjusted to be independent of text length.
9 Causation terms: Terms such as because, *effect*, that try to assign a cause to whatever the person is describing.
10 Certainty terms: Words such as *always*, *never* that are used to add concreteness or definitiveness to the language.
11 Cognitive processing terms: Words related to mental processes such as cause and *know*.
12 Content word diversity: The number of unique content words that are used.
13 Emotiveness: Ratio of adjectives and adverbs to nouns and verbs.
14 Exclusive terms: Terms such as *without*, *not*, *but*, *exclude*.
15 Generalizing terms: Terms such as *everyone*, *all*, *anybody*.
16 Imagery: Extent to which language is easy or hard to imagine.
17 Lexical diversity: The number of unique words that are used. Unlike bilogarithmic type-token ratio, this measure is not adjusted to be independent of text length.
18 Modal verbs: Auxiliary verbs such as *can*, *must*, *shall*, that provide additional information about the related verb.
19 Modifiers: Adjectives and adverbs.
20 Motion terms: Terms such as arrive, go that describe movement.
21 Passive verbs: Uses of the "to be" verbs such as *is*, *was*.
22 Pausality: Number of punctuation marks/number of sentences.
23 Pleasantness: Degree to which word is pleasant or unpleasant.
24 Redundancy: Ratio of function words to number of sentences. Function words, such as articles and pronouns, are used to form grammatical relationships between other words.
25 Sensory ratio: Ratio of words related to sensation, such as *taste*, *touch*, *feel*, *hear*.
26 Sentence quantity: The number of sentences in the text.
27 Spatial ratio: Words related to locations of people or objects, such as *inside* or *under*.
28 Temporal ratio: Words related to time or timing such as *yesterday*, *tomorrow*, *later*, *frequently*.
29 Tentative terms: Words such *maybe*, *perhaps*, *seem*.
30 Verb quantity: The number of verbs that are used.
31 Word quantity: The total number of words that are used.

Output feature (the dependent variable):

1 Deception?: False, the statement was deceptive, or True, the statement truthful.

Feature descriptions are partially adapted from Pennebaker and Francis (2001), Zhou, Burgoon, Nunamaker et al. (2004), Zhou, Burgoon, Twitchell et al. (2004).

types of monotone and/or non-monotone relationships, this value can be changed). Categorical predictors are not transformed in any way. This feature selection algorithm does not assume any particular type or shape of relationship between the predictors and the dependent variables (classes) of interest. Instead, it employs a generalized "notion of relationship" while screening the predictors, one by one, for the classification problems. Fig. 2 illustrates the results of the *Chi-square* statistics-based feature selection algorithm. From the 31 cues, the top 13 features are included in the model development efforts. The considerable drop in Chi-square value after the 13th feature suggests those that remain may not produce sufficient return when weighted against their complexity impact on the model.

3.3. Detection models

In this study, we used three data mining methods (artificial neural networks, decision trees and logistic regression) along with an information fusion-based ensemble method. The first three are arguably the most popular methods used in a wide range of applied data mining studies, and hence are chosen herein to provide a baseline for the accuracy as well as a comparison point to previously published studies. What follows is a brief description of these detection models.

3.3.1. Artificial neural networks

Artificial neural networks (ANN) are commonly known as biologically inspired analytical techniques, capable of modeling extremely complex non-linear functions (Haykin, 2008). In this study we used a popular neural network architecture called Multi-Layer Perceptron (MLP) with a back-propagation learning algorithm. MLP is essentially the collection of nonlinear neurons organized and connected to each other in a feed-forward multi-layered structure.

3.3.2. Decision trees

Decision trees, as the name implies, is a technique that recursively separates observations in branches to construct a tree for the purpose of improving the detection accuracy (Breiman, Friedman, Olshen, & Stone, 1984). In doing so, different mathematical algorithms (e.g., entropy-based information gain, Gini index, etc.) are use to identify a variable and the corresponding threshold for the variable that splits the pool of observations into two or more subgroups. This step is repeated at each leaf node until the complete tree is constructed. The specific decision tree model used here was C&RT.

3.3.3. Logistic regression

Logistic regression is a generalization of linear regression. It is used primarily for predicting binary or multi-class dependent variables. Because the response variable is discrete, it cannot be modeled directly by linear regression. Therefore, rather than predicting a point estimate of the event itself, it builds the model to predict the odds of its occurrence. In a two-class problem, odds greater than 50% means that the case is assigned to the class designated as "1" and "0" otherwise. While logistic regression is a very powerful modeling tool, it assumes that the response variable (the log odds, not the event itself) is linear in the coefficients of the predictor variables. Furthermore, the modeler, based on his or her experience with the data and data analysis, must choose the right inputs and specify their functional relationship to the response variable.

3.3.4. Information fusion

Information fusion-based ensemble methods use a process of "intelligently" combining the information (detections in this case) provided by two or more information sources (i.e., detection models). For example, if a statement is labeled deceptive by two or more of the three individual models used here, it is classified as deceptive, while a statement classified as deceptive by only one of the models is labeled as truthful. While there is an ongoing debate about the sophistication level of the fusion methods, there is a general consensus that fusion (combining detections) usually produces more accurate and more robust detection models.

4. Results

Table 2 shows the results of the three data mining methods as well as the results of the information fusion-based ensemble method. These results are obtained using a 10-fold cross validation methodology; that is for each model type, 10 different prediction models are built and tested using a mutually exclusive 10% sample of the total dataset. In Table 2, the first two rows show the confusion matrixes (i.e., coincidence matrix or classification matrix constructed from the test data samples) for all prediction model types. In the confusion matrixes, the columns represent the actual classes while the rows represent the mode predictions. The third row shows the detection accuracy for each class (deceptive [F] and truthful [T]) separately while the fourth row shows the overall detection accuracy for each model type. As the results indicate, for the individual models, ANN performed the best with 73.46% accuracy on test data samples while the decision tree model performed the second best with 71.60% accuracy. The ensemble model (which is the linear combination of the other three algorithm detections) performed slightly better in detection accuracy than ANN with 74.07% accuracy. Besides the fact that even a small difference in accuracy can become very important in a high-stakes environment, this type of ensemble models is preferred because they are based on the combination of multiple model outcomes and hence are more robust and reliable.

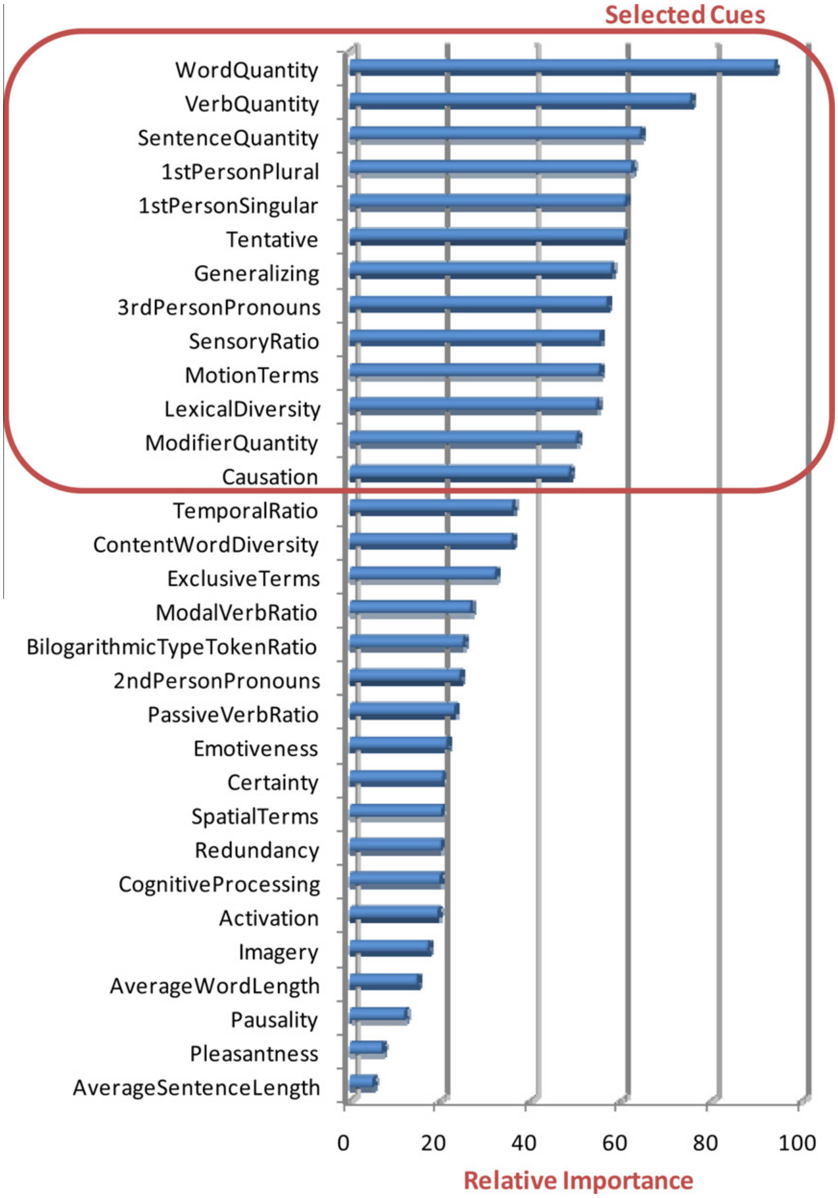

Fig. 2. Chi-square-based feature selection results.

Table 2
Detection results on test data sample for all models.

		ANN		C&RT		Logit		Ensemble	
		T	F	T	F	T	F	T	F
	T	62	24	57	22	55	27	61	22
	F	19	57	24	59	26	54	20	59
Class accuracy		76.54	70.37	70.37	72.84	67.90	66.67	75.31	72.84
Overall Acc. (%)		73.46		71.60		67.28		74.07	

5. Discussion and conclusion

These results show that automated text-based deception detection has the potential to aid those who must try to detect lies in text. We have shown that the combination of text and data mining techniques can be successfully applied to real-world data. Here, a military crime scenario was investigated, though the technique has the potential to be implemented in a wide variety of fields. The accuracy here (approximately 74%) has exceeded the 63% accuracy of the PCASS mock crime studies, an important first step in showing that this technique can be applied to real-world problems. However, the accuracy of the tool must be tested across different contexts and cultures to know how widely these results will apply. Despite these promising results, the study was not without complications, as described below.

Though we have accurately tested the technique for a specific domain, this tool will need to be calibrated with samples from additional contexts. While some small set of features may carry across domains, results to date suggest that the pertinent cues will vary in different circumstances. For example, language use in severe, or high-stakes, circumstances may differ from that in more ordinary situations. Until or unless a universal set of features is identified, one approach might be to always include all possible features. Alternatively, many feature selection procedures might be attempted to maximize accuracy. Additional relevant features may emerge in the literature as deception research continues. These features could then be integrated into the system. Further, while some of the features may be well-supported in deception literature, the list of terms (or dictionary) that is associated with some of the features may need improvement.

Automatic text-based deception detection has focused on the combination of data mining algorithms and content-independent linguistic cues that can be extracted using part of speech taggers or lists of terms. Currently, these cues are derived from deception theory. By using well-understood data mining techniques along with clearly defined cues, replication across contexts and samples is facilitated, providing a baseline for comparison for future efforts. As the accuracy of these techniques peaks, alternative text-mining technologies may need to be pursued to gain further advancements. In this case, the limited sample sizes, as well as the need for portable, user-friendly methods must continue to be incorporated in the deception detection task.

One of the biggest challenges of this study was gathering a data set of sufficient size. While many person-of-interest statements are recorded each year, only a limited number could be conclusively identified as truthful or deceptive. Sample size is a recurring issue in deception research. Therefore, data mining techniques which are designed for large sample sizes must be carefully adapted to these smaller samples. That is, while mining smaller datasets (as was the case in this study), careful selection and proper use of variable reduction (e.g., Chi-square test) and experimental design (e.g., 10-fold cross validation) methods become crucial.

There remains some work to do before this technique can be widely implemented. As previously mentioned, representative data samples will need to be collected so that the system can be calibrated for different domains. Currently, the text-processing and classification are conducted separately. A user interface will need to be developed and tested so these are combined into a seamless process. This interface will allow users to input data to be evaluated. Finally, appropriate output reports will need to be developed to enable the user to easily evaluate whether deception took place. This work can be accomplished with available technology and can feasibly be implemented in a handheld device. When fully developed, automated text-based deception detection may help fill the need for an accurate, user-friendly, non-invasive, portable credibility assessment device.

Acknowledgements

This project was partially funded by the Defense Academy for Credibility Assessment, Project DoDPI06-P-0041, Contract W991NF-07-1-0375. The views expressed in this article are those of the authors and do not necessarily reflect the official policy or position of the Department of Defense or the U.S. Government.

References

Adkins, M., Twitchell, D., Burgoon, J. K., & Nunamaker Jr, J. F. (2004). Advances in automated deception detection in text-based computer-mediated communication. In: *Paper presented at the enabling technologies for simulation science VIII, Orlando, FL, USA*.

Battelle Memorial Institute (2007). Efficacy of prototype credibility assessment technologies: PCASS final report.

Berry, M. J. A., & Linoff, G. S. (2004). *Data mining techniques* (2nd ed.). Indianapolis, Indiana: Wiley Publishing.

Bond, C. F., & DePaulo, B. M. (2006). Accuracy of deception judgments. *Personality and Social Psychology Reports, 10*(3), 214–234.

Bond, G. D., & Lee, A. Y. (2005). Language of lies in prison: Linguistic classification of prisoners' truthful and deceptive natural language. *Applied Cognitive Psychology, 19*(3), 313.

Breiman, L., Friedman, J. H., Olshen, R. A., & Stone, C. J. (1984). *Classification and regression trees*. Monterey, CA: Wadsworth & Brooks Books & Software.

Buller, D. B., & Burgoon, J. K. (1996). Interpersonal deception theory. *Communication Theory, 6*(3), 203–242.

Cunningham, H. (2002). Gate, a general architecture for text engineering. *Computers and the Humanities, 36*(2), 223–254.

DePaulo, B. M., Lindsay, J. J., Malone, B. E., Muhlenbruck, L., Charlton, K., & Cooper, H. (2003). Cues to deception. *Psychological Bulletin, 129*(1), 74–118.

Feldman, R., & Sanger, J. (2007). *The text mining handbook: Advanced approaches in analyzing unstructured data*. New York: Cambridge University Press.

Frank, M. G., & Feeley, T. H. (2003). To catch a liar: Challenges for research in lie detection training. *Journal of Applied Communication Research, 31*(1), 58–75.

George, J. F., & Keane, B. T. (2006). Deception detection by third party observers. In: *Paper presented at the deception detection symposium, 39th annual Hawaii international conference on system sciences*.

Hancock, J., Thom-Santelli, J., & Ritchie, T. (2004). Deception and design: The impact of communication technology on lying behavior. In: *Paper presented at the SIGCHI conference on Human factors in computing systems, Vienna, Austria, April 24–29*.

Hancock, J. T., Curry, L., Goorha, S., & Woodworth, M. (2005). Automated linguistic analysis of deceptive and truthful synchronous computer-mediated communication. In: *Paper presented at the 38th annual Hawaii international conference on system sciences*.

Harris, J. C., & McQuarrie, A. D. (2009). The preliminary credibility assessment system embedded algorithm description and validation results. Johns Hopkins University Applied Physics Laboratory Report Number GED-R-06-7571.

Haykin, S. (2008). *Neural networks and learning machines* (3rd ed.). New Jersey: Prentice Hall.

Indictment puts internet pranksters on notice (2008). An article on CNN Website. <http://www.cnn.com/2008/CRIME/05/16/internet.suicide.ap/index.html> Accessed 16.5.08.

Johnson, M. K., & Raye, C. L. (1981). Reality monitoring. *Psychological Review, 88*(1), 67–85.

Kohnken, G. (1985). Speech and deception of eyewitnesses: An information processing approach. In F. L. Denmark (Ed.), *Social/ecological psychology and the psychology of women*. North-Holland: Elsevier Science Publishers.

McCornack, S. A. (1992). Information manipulation theory. *Communication Monographs, 59*(1), 1–16.

National Research Council (2003). The polygraph and lie detection. Technical report, Washington, DC.

Pennebaker, J. W., & Francis, M. E. (2001). *Linguistic inquiry and word count: Liwc 2001*. Mahwah, NJ: Erlbaum Publishers.

Pollina, D. A., Dollins, A. B., Senter, S. M., Krapohl, D. J., & Ryan, A. H. (2004). Comparison of polygraph data obtained from individuals involved in mock crimes and actual criminal investigations. *Journal of Applied Psychology, 89*(6), 1099–1105.

Senter, S. M., Waller, J., & Krapohl, D. J. (2006). *Validation studies for the preliminary credibility assessment screening system (pcass)*. Department of Defense Polygraph Institute.

Twitchell, D., Jensen, M. L., Burgoon, J. K., & Nunamaker Jr., J. F., (2004). Detecting deception in secondary screening interviews using linguistic analysis. In: *Paper presented at the 7th international IEEE conference on intelligent transportation systems*.

Twitchell, D. P., Biros, D. P., Adkins, M., Forsgren, N., Burgoon, J. K., & Nunamaker Jr, J. F. (2006). Automated determination of the veracity of interview statements from people of interest to an operational security force. In: *Paper presented at the 39th annual Hawaii international conference on system sciences*.

Vrij, A. (2000). *Detecting lies and deceit: The psychology of lying and the implications for professional practice*. New York: John Wiley & Sons.

Zhou, L., Burgoon, J. K., Nunamaker, J., Jay, F., & Twitchell, D. P. (2004). Automated linguistics based cues for detecting deception in text-based asynchronous computer-mediated communication: An empirical investigation. *Group Decision and Negotiation, 13*(1), 81–106.

Zhou, L., Burgoon, J. K., Twitchell, D. P., Qin, T. T., & Nunamaker, J. F. (2004). A comparison of classification methods for predicting deception in computer-mediated communication. *Journal of Management Information Systems, 20*(4), 139–165.