

Identification of Truth and Deception in Text: Application of Vector Space Model to Rhetorical Structure Theory

Victoria L. Rubin and Tatiana Vashchilko

Language and Information Technology Research Lab (LiT.RL)
Faculty of Information and Media Studies, University of Western Ontario
London, Ontario, Canada
{vrubin, tvashchi}@uwo.ca

Abstract

The paper proposes to use Rhetorical Structure Theory (RST) analytic framework to identify systematic differences between deceptive and truthful stories in terms of their coherence and structure. A sample of 36 elicited personal stories, self-ranked as completely truthful or completely deceptive, is manually analyzed by assigning RST discourse relations among a story's constituent parts. Vector Space Model (VSM) assesses each story's position in multi-dimensional RST space with respect to its distance to truth and deceptive centers as measures of the story's level of deception and truthfulness. Ten human judges evaluate if each story is deceptive or not, and assign their confidence levels, which produce measures of the human expected deception and truthfulness levels. The paper contributes to deception detection research and RST twofold: a) demonstration of discourse structure analysis in pragmatics as a prominent way of automated deception detection and, as such, an effective complement to lexico-semantic analysis, and b) development of RST-VSM methodology to interpret RST analysis in identification of previously unseen deceptive texts.

Introduction

Automated deception detection is a challenging task (DePaulo, Charlton, Cooper, Lindsay, and Muhlenbruck, 1997), only recently proven feasible with natural language processing and machine learning techniques (Bachenko, Fitzpatrick, and Schonwetter, 2008; Fuller, Biros, and Wilson, 2009; Hancock, Curry, Goorha, and

Woodworth, 2008; Rubin, 2010; Zhou, Burgoon, Nunamaker, and Twitchell, 2004). The idea is to distinguish truthful information from deceptive, where deception usually implies an intentional and knowing attempt on the part of the sender to create a false belief or false conclusion in the mind of the receiver of the information (e.g., Buller and Burgoon, 1996; Zhou, et al., 2004). In this paper we focus solely on textual information, in particular, in computer-mediated personal communications such as e-mails or online posts.

Previously suggested techniques for detecting deception in text reach modest accuracy rates at the level of lexico-semantic analysis. Certain lexical items are considered to be predictive linguistic cues, and could be derived, for examples, from the Statement Validity Analysis techniques used in law enforcement for credibility assessments (as in Porter and Yuille, 1996). Though there is no clear consensus on reliable predictors of deception, deceptive cues are identified in texts, extracted and clustered conceptually, for instance, to represent diversity, complexity, specificity, and non-immediacy of the analyzed texts (e.g., Zhou, Burgoon, Nunamaker, and Twitchell (2004)). When implemented with standard classification algorithms (such as neural nets, decision trees, and logistic regression), such methods achieve 74% accuracy (Fuller, et al., 2009). Existing psycholinguistic lexicons (e.g., LWIC by Pennebaker and Francis, 1999) have been adapted to perform binary text classifications for truthful versus deceptive opinions, with an average classifier demonstrating 70% accuracy rate (Mihalcea and Strapparava, 2009).

These modest results, though usually achieved on restricted topics, are promising since they supersede notoriously unreliable human abilities in lie-truth discrimination tasks. On average, people are not very good at spotting lies (Vrij, 2000), succeeding generally only about half of the time (Frank, Paolantini, Feeley, and

Servoss, 2004). For instance, a meta-analytical review of over 100 experiments with over 1,000 participants, showed a 54% mean accuracy rate at identifying deception (DePaulo, et al., 1997). Human judges achieve 50 – 63% success rates, depending on what is considered deceptive on a seven-point scale of truth-to-deception continuum (Rubin and Conroy, 2011, Rubin and Conroy, 2012), but the higher the actual self-reported deception level of the story, the more likely a story would be confidently assigned as deceptive. In other words, extreme degrees of deception are more transparent to judges.

The task for current automated deception detection techniques has been formulated as binary text categorization – is a message deceptive or truthful – and the decision applies to the whole analyzed text. Since it is an overall discourse level decision, it may be reasonable to consider discourse or pragmatic features of each message. Thus far, discourse is surprisingly rarely considered, if at all, and the majority of the effort has been restricted to lexico-semantic verbal predictors. A rare exception up to date has been a Bachenko, Fitzpatrick and Schonwetter's (2008) study that focuses on truth or falsity of individual propositions, achieving a finer-grained level of analysis¹, but the propositional inter-relations within the discourse structure are not considered. To the best of our knowledge there have been no advances in that automation deception detection task to incorporate discourse structure features and/or text coherence analysis at the pragmatic levels of story interpretation.

Study Objective

With the recent advances in the identification of verbal cues of deception in mind, and the realization that they focus on linguistic levels below discourse and pragmatic analysis, the study focuses on one main question:

- What is the impact of the relations between discourse constituent parts on the discourse composition of deceptive and truthful messages?

We hypothesize that if the relations between discourse constituent parts in deceptive messages differ from the ones in truthful messages, then systematic analysis of such relations will help to

¹ Using a corpus of criminal statements, police interrogations and legal testimonies, their regression and tree-based classification automatic tagger performs at average 69% recall and 85% precision rates, as compared to the performance of human taggers on the same subset (Bachenko, et al., 2008).

detect deception. To investigate this question, we propose to use a novel methodology for deception detection research, Rhetorical Structure Theory (RST) analysis with subsequent application of the Vector Space Model (VSM). RST analysis is promising in deception detection, since RST analysis captures coherence of a story in terms of functional relations among different meaningful text units, and describes a hierarchical structure of each story (Mann and Thompson, 1988). The result is that each story is a set of RST relations connected in a hierarchical manner with more salient text units heading this hierarchical tree. We also propose to utilize the VSM model for conversion of the derived RST relations' frequencies into meaningful clusters of diverse deception levels. To evaluate the proposed RST-VSM methodology of deception detection in texts, we compare human assessment to the RST-analysis of deception levels for the sets of deceptive and truthful stories. The main findings demonstrate that RST resembles, to some degree, human judges in deceptive and truthful stories, and RST deception detection in self-rated deceptive stories has greater consistency than in truthful ones, which signifies the prominence of using RST-VSM methodology for deception detection². However, RST conclusions regarding levels of deception in the truthful stories requires further research about the diversity of RST relations for the expressions of truths and deception as well as the types of clustering algorithms most suitable for clustering unevaluated by human judges' written communication in RST space to detect deception with certain degree of precision.

The paper has three main parts. The next part discusses methodological foundations of RST-VSM approach. Then, the data and collection method describe the sample. Finally, the results section demonstrates the identified levels of deception and truthfulness as well as their distribution across truthful and deceptive stories.

RST-VSM Methodology: Combining Vector Space Model and Rhetorical Structure Theory

Vector space model (VSM) seemed to be very useful in the identification of truth and deception types of written stories especially if the meaning

² The authors recognize that the results are preliminary and should be generalized with caution due to very small dataset and certain methodological issues that require further development.

of the stories is represented as RST relations. RST differentiates between rhetorically stand-alone parts of a text, some of which are more salient (nucleolus) than the others (satellite). In the past couple of decades, empirical observations and previous RST research confirmed that writers tend to emphasize certain parts of a text in order to express their most essential idea to reach the purpose of the written message. These parts can be systematically identified through the analysis of the rhetorical connections among more and less essential parts of a text. RST helps to describe and quantify text coherence through a set of constraints on nucleolus and satellites. The main function of these constraints is to describe in the meaningful way why and how one part of a text connects to the others within a hierarchical tree structure, which is an RST representation of a coded text. The names of the RST relations also resemble the purpose of using the connected text parts together.

For example, one of the RST relations, which appear in truthful stories and never appear in the deceptive stories in our sample, is EVIDENCE. The main purpose of using EVIDENCE to connect two parts of text is to present additional information in satellite, so that the reader's belief about the information in the nucleolus increases. However, this can happen only if the information in the satellite is credible from reader's point of view. For some reason, the RST coding of 18 deceptive stories has never used EVIDENCE, but used it rather often in 18 truthful stories. This might indicate that either 1) writers of deceptive stories did not see any purpose in supplying additional information to the readers to increase their beliefs in communicating writer's essential ideas, or 2) the credibility of presented information in satellite was not credible from the readers' points of view, which did not qualify the relationship between nucleolus and satellite for "EVIDENCE" relation, or 3) both (See an example of RST diagram in Appendix A).

Our premise is that if there are systematic differences between deceptive and truthful written stories in terms of their coherence and structure, then the RST analysis of these stories can identify two sets of RST relations and their structure. One set is specific for the deceptive stories, and the other one is specific for the truthful stories.

We propose to use a vector space model for the identification of these sets of RST relations. Mathematically speaking, written stories have to

be modeled in a way suitable for the application of various computational algorithms based on linear algebra. Using a vector space model, the written stories could be represented as RST vectors in a high dimensional space (Salton and McGill 1983, Manning and Schutse 1999). According to the VSM, stories are represented as vectors, and the dimension of the vector space equals to the number of RST relations in a set of all written stories under consideration. Such representation of written stories makes the VSM very attractive in terms of its simplicity and applicability (Baeza-Yates and Ribeiro-Neto 1999).

Vector space model³ is the basis for almost all clustering techniques when dealing with the analysis of texts. Once the texts are represented according to VSM, as vectors in an n-dimensional space, we can apply the myriad of cluster methods that have been developed in Computational Science, Data Mining, Bioinformatics. Cluster analysis methods can be divided into two big groups (Zhong and Ghosh 2004): discriminative (or similarity based) approaches (Indyk 1999, Scholkopf and Smola 2001, Vapnik 1998) and generative (or model-based) approaches (Blimes 1998, Rose 1998, Cadez et al. 2000).

The main benefit of applying vector space model to RST analysis is that the VSM allows a formal identification of coherence and structural similarities among stories of the same type (truthful or deceptive). For this purpose, RST relations are vectors in a story space. Visually we could think about the set of stories or RST relations as a cube (Figure 1), in which each dimension is an RST relation.




Figure 1: Cluster Representation of Story Sets or RST Relations (Cluto Graphical Frontend Project, 2002).

³ Tombros (2002) maintains that most of the research related to the retrieval of information is based on vector space model.

The main subsets of this set of stories are two clusters, deceptive stories and truthful stories. The element of a cluster is a story, and a cluster is a set of elements that share enough similarity to be grouped together, the deceptive stories or truthful stories (Berkhin 2002). That is, there is a number of distinctive features (RST relations, their co-occurrences and positions in a hierarchical structure) that make each story unique and being a member of a particular cluster. These distinctive features of the stories are compared, and when some similarity threshold is met, they are placed in one of two groups, deceptive or truthful stories.

Similarity⁴ is one of the key concepts in cluster analysis, since most of the classical techniques (k-means, unsupervised Bayes, hierarchical agglomerative clustering) and rather recent ones (CLARANS, DBSCAN, BIRCH, CLIQUE, CURE, etc.) “are based on distances between the samples in the original vector space” (Strehl et al 2000). Such algorithms form a similarity based clustering framework (Figure 1) as it is described in Strehl et al (2000) , or as Zhong and Ghosh (2004) define it as discriminative (or similarity – based) clustering approaches.

That is why, this paper modifies Strehl et al’s (2004) similarity based clustering framework (Figure 2) to develop a unique RST-VSM methodology for deception detection in text. The RST-VSM methodology includes three main steps:

1) The set of written stories, X , is transformed into the vector space description, X , using some rule, Y , that in our case corresponds to an RST analysis and identification of RST relations as well as their hierarchy in each story.

2) This vector space description X is transformed into a similarity space description, S , using some rule, Ψ , which in our case is the Euclidian distance of every story from a deception and truth centers correspondingly based on normalized frequency of RST relations in a written story⁵.

3) The similarity space description, S , is mapped into clusters based on the rule Φ , which we define as the relative closeness of a story to a

deception or a truth center: if a story is closer to the truth center, then a story is placed in a truth cluster, whereas if a story is closer to a deception center, then a story is placed in a deception cluster.




Figure 2: Similarity Based Clustering Framework (Strehl et al, 2004)

Data Collection and Sample

The dataset contains 36 rich unique personal stories, elicited using Amazon’s online survey service, Mechanical Turk (www.mturk.com). Respondents in one group were asked to write a rich unique story, which is completely true or with some degree of deception. Respondents in another group were asked to evaluate the stories written by the respondents in the first group (For further details on the data collection process and the discussion of associated challenges, see Rubin and Conroy 2012).

Two groups of 18 stories each compile the data sample. The first group consists of 18 stories that were self-ranked by their authors as completely deceptive on a seven-point Likert scale from complete truth to complete deception (deceptive self-reported group). The second group includes stories, which their authors rated as completely truthful stories (truthful self-reported group). The second group was matched in numbers for direct comparisons to the first group by selecting random 18 stories from a larger group of 39 completely truthful stories (Rubin and Conroy, 2011, Rubin and Conroy, 2012). Each story in both groups, truthful self-reported and deceptive self-reported, has 10 unique human judgments associated with it. Each judgment is binary (“judged truthful” or “judged deceptive”), and has an associated confidence level assigned by the judge (either “totally uncertain”, “somewhat uncertain”, “I’m guessing”, “somewhat certain”, or “totally certain”). Each writer and judge was encouraged to provide explanations for defining a story as truthful or deceptive, and assigning a particular confidence level. In total, 396 participants contributed to the study, 36 of them were story authors, and 360 – were judges performing lie-truth discrimination task by confidence level.

⁴ “Interobject similarity is a measure of correspondence or resemblance between objects to be clustered” (Hair et al. 1995, p. 429).

⁵ Since RST stories as vectors differ in length, the normalization assures their comparability. The coordinates of every story (the frequency of an RST relation in a story) are divided on the vector’s length.

We combine the 10 judges' evaluations of a story into one measure, the expected level of a story's deception or truthfulness. Since judges' confidence levels reflect the likelihood of a story being truthful or deceptive, the probability of a story being completely true or deceptive equals one and corresponds to a "totally certain" confidence level that the story is true or deceptive⁶. Two dummy variables are created for each story. One dummy, a deception dummy, equals 1, if a judge rated the story is "judged deceptive", and 0 otherwise. The second dummy, the truthfulness dummy, equals 1 if a judge rated the story as "judged truthful", and 0 otherwise. Then the expected level of deception of a story equals the product of the probability (confidence level) of deception and the deception dummy across 10 judges. Similarly, the expected level of truthfulness is equals the product of the probability of truthfulness (confidence level) and the truthfulness dummy across 10 judges. The distribution of expected levels of deception and the expected levels of truthfulness of the deceptive and truthful subsets of the sample are in Appendix B1-B2.

Thirty six stories, evenly divided between truthful and deceptive self-report groups, were manually analyzed using the classical set of Mann and Thompson's (1988) RST relations, extensively tested empirically (Taboada and Mann, 2006). As a first stage of RST-VSM methodology development, the manual RST coding was required to deepen the understanding of the rhetorical relations and structures specific for deceptive and truthful stories. Moreover, manual analysis aided by Mick O'Donnell's RSTTool (<http://www.wagsoft.com/RSTTool/>) might ensure higher reliability of the analysis and avoid compilation of errors, as the RST output further served as the VSM input. Taboada (2004) reports on the existence of Daniel Marcu's RST Annotation Tool: www.isi.edu/licensed-sw/RSTTool/ and Hatem Ghorbel's RhetAnnotate (http://www.epfl.ch/~ghorbel/rhet_annotate/) and provides a good overview of other recent RST resources and applications. The acquired knowledge during manual coding of deceptive stories along with recent advances in automated RST analysis will help later on to evaluate RST-VSM methodology and design a

⁶ In the same way, the other levels of confidence have the following probability correspondences: "totally uncertain" has probability 0.2 of a story being deceptive or truthful, "somewhat uncertain" – 0.4, "I'm guessing" – 0.6, and "somewhat certain" – 0.8.

completely automated deception detection tool relying on the automated procedures to recognize rhetorical relations, which utilize the full rhetorical parsing (Marcu 1997, 2002).

Results

The preliminary clustering of 36 stories in RST space using various clustering algorithms shows that RST dimensions can systematically differentiate between truthful and deceptive stories as well as diverse levels of deception (Figure 3).

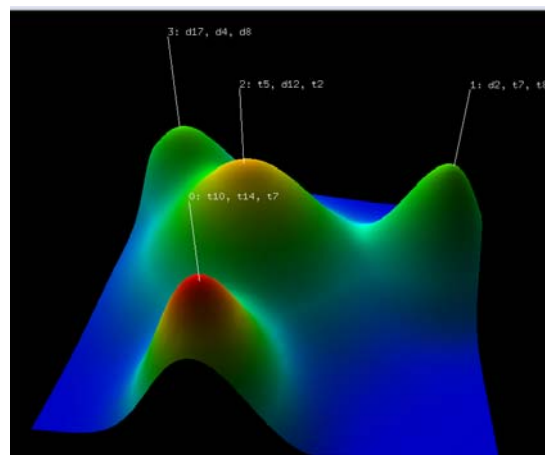


Figure 3. Four Clusters in RST Space by Level of Deception.

The visualization uses GLUTO software (<http://glaros.dtc.umn.edu/gkhome/cluto/gcluto/overview>), which finds the clustering solution as a result of the optimization of a "particular function that reflects the underlying definition of the "goodness" of the cluster" (Rasmussen and Karypis 2004, p.3). Among the four clusters in RST space, two clusters are composed of completely deceptive stories (far back left peak in green) or entirely truthful stories (front peak in red), the other two clusters have a mixture with the prevalence of either truthful or deceptive stories. This preliminary investigation of using RST space for deception detection indicates that the RST analysis seems to offer a systematical way of distinguishing between truth and deceptive features of texts.

This paper develops an RST-VSM methodology by using RST analysis of each story in N-dimensional RST space with subsequent application of vector space model to identify the level of a story's deception. A normalized frequency of an RST relation in a story is a distinct coordinate in the RST space. The authors' ratings are used to calculate the

centers for the truth and deception clusters based on corresponding authors' self-rated deception and truthful sets of stories in the sample. The normalized Euclidian distances between a story and each of the centers are defined as the degree of deception of that story depending on its closeness to the deception center. The closer a story is to the deception center, the higher is its level of deception. The closer a story is to the truthful center, the higher is its level of truthfulness⁷.

RST seems to differentiate between truthful and deceptive stories. The difference in means test demonstrates that the truthful stories have a statistically significantly lower average number of text units per statement than the deceptive stories ($t = -1.3104$), though these differences are not large, only at 10% significance level. The normalized frequencies of the RST relations appearing in the truthful and deceptive stories differ for about one third of all RST relations based on the difference in means test (Appendix C).

The comparison of the distribution of RST relations across deceptive and truth centers demonstrates that on average, the frequencies and the usage of such RST relations as conjunction, elaboration, evaluation, list, means, non-volitional cause, non-volitional result, sequence, and solutionhood in deceptive stories exceeds those in the truthful ones (Figure 4). On the other hand, the average usage and frequencies of such RST relations as volitional result, volitional cause, purpose, interpretation, concession, circumstance and antithesis in truthful stories exceeds those in the deceptive ones. Some of the RST relations are only specific for one type of the story: enablement, restatement and evidence appear only in truthful stories, whereas summary, preparation, unconditional and disjunction appear only in deceptive stories.

The histograms of distributions of deception (truthfulness) levels assigned by judges and derived from RST-VSM analysis demonstrate some similarities between the two for truth and for deceptive stories (Appendices D-E). More rigorous statistical testing reveals that only truthfulness levels in deceptive stories assigned by judges do not have statistically significant difference from the RST-VSM ones⁸. For other

groups, the judges' assessments and RST ones do differ significantly.




Figure 4: Comparison of the RST Relations' Composing the Deceptive Cluster Center (top red bar) and the Truthful Cluster Center (bottom blue bar).

The difference is especially apparent in the distributions of deception and truthfulness in truthful stories. Among them, RST-VSM methodology counted 44.44% of stories having 50% deception level, whereas judges counted 61.11 percent of the same stories having low deception level of no more than 20%. The level of truthfulness was also much higher in judges' assessment than based on RST-VSM calculations.

⁷ All calculations are performed in STATA.

⁸ We use the Wilcoxon signed rank sum test, which is the non-parametric version of a paired samples t-test (STATA command `signrank` (STATA 2012)).

The distribution of the levels of deception and truthfulness across all deceptive stories (Appendices D1-D4) and across all truthful stories (Appendices E1-E4) shows variations in patterns of deception levels based on RST-VSM. In deception stories, the RST-VSM levels of deception are consistently higher than the RST-VSM levels of truthfulness. Assuming that the authors of the stories did make them up, the RST-VSM methodology seems to offer a systematic way of detecting a high level of deception with rather good precision.

The RST-VSM deception levels are not as high as human judges' ones, with human judges assigning much higher levels of deception to deceptive stories than to truthful stories. Assuming that the stories are indeed made up, the human judges have greater precision than the RST-VSM methodology. Nevertheless, RST-VSM analysis assigns higher deception levels to stories, which also receive higher human judges' deception levels. This pattern is consistent across all deceptive stories.

Discussion

The analysis of truthful stories shows some systematic and some slightly contradictory findings. On one hand, the levels of truthfulness assigned by judges are predominantly higher than the levels of deception. Again, assuming that the stories in the truthful set are completely true because the authors ranked them so, the human judges have greater likelihood of rating these stories as truthful than as deceptive. This can be an indicator of a good precision of deception detection by human judges.

On the other hand, the RST-VSM analysis also demonstrates that large subsample (but not as large as indicated by human judges) of truthful stories is closer to the truth center than to the deceptive one. However, it seems that RST-VSM methodology overestimates the levels of deception in the truthful stories compared to human judges.

Overall, however, the RST-VSM analysis demonstrates a positive support for the proposed hypothesis. The apparent and consistent closeness of deceptive stories to RST deception center (compared to the closeness of the deceptive stories to the truthful center) and truthful stories to RST truthful center can indicate that the relations between discourse constituent parts differ between truthful and deceptive messages. Thus, since the truthful and

deceptive relations exhibit systematic differences in RST space, the proposed RST-VSM methodology seemed to be a prominent tool in deception detection. The results, however, have to be interpreted with caution, since the sample was very small, and only one expert conducted RST coding.

The discussion, however, might be extended to the case, where the assumption of self-ranked levels of deception and truthfulness do not hold. In this case we still suspect that even deceptive story might contain elements of truth (though much less), and the truth story will have some elements of deception. RST-VSM analysis demonstrated greater levels of deception in truth and deceptive stories compared to the human judges. This might indicate that RST-VSM potentially offers an alternative to human judges way of detecting deception when it is least expected in text (as in the example of supposedly truthful stories) or detecting it in a more accurate way (if some level of deception is assumed as in the completely deceptive stories). The advantage of RST-VSM methodology is in its rigorous and systematic approach of coding discourse relations and their subsequent analysis in RST space using vector space models. As a result, the relations between units exhibiting different degrees of salience in text because of writers' purposes with their subsequent readers' perceptions become indicators of diversity in deception levels.

Conclusions

To conclude, relations between discourse parts along with its structure seem to have different patterns in truthful and deception stories. If so, RST-VSM methodology can be a prominent way of detecting deception and complementing the existing lexical ones.

Our contribution to deception detection research and RST twofold: a) we demonstrate that discourse structure analysis and pragmatics as a promising way of automated deception detection and, as such, an effective complement to lexico-semantic analysis, and b) we develop the unique RST-VSM methodology of interpreting RST analysis in identification of previously unseen deceptive texts.


Acknowledgments

This research is funded by the New Research and Scholarly Initiative Award (10-303) from the Academic Development Fund at Western.

References


- Bachenko, J., Fitzpatrick, E., and Schonwetter, M. 2008. Verification and implementation of language-based deception indicators in civil and criminal narratives. In *Proceedings of the 22nd International Conf. on Computational Linguistics*.
- Baeza-Yates, R. and B. Ribeiro-Neto. 1999. *Modern Information Retrieval*. New York: Addison-Wesley
- Buller, D. B., and Burgoon, J. K. 1996. Interpersonal Deception Theory. *Communication Theory*, 6(3), 203-242.
- Berkhin, P. 2002. Survey of Clustering Data Mining Techniques. DOI: 10.1.1.18.3739.
- Blimes, J. A. 1998. *A Gentle Tutorial of the EM Algorithm and Its Application to Parameter Estimation for Gaussian Mixture and Hidden Markov Models*: Univ. of California, Berkeley.
- Cadez, I. V., Gaffney, S. and P. Smyth. 2000. A General Probabilistic Framework for Clustering Individuals and Objects. In *Proceedings of the 6th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- DePaulo, B. M., Charlton, K., Cooper, H., Lindsay, J. J., and Muhlenbruck, L. 1997. The Accuracy-Confidence Correlation in the Detection of Deception. *Personality and Social Psychology Review*, 1(4), 346-357.
- Frank, M. G., Paolantino, N., Feeley, T., and Servoss, T. 2004. Individual and Small Group Accuracy in Judging Truthful and Deceptive Communication. *Group Decision and Negotiation*, 13, 45-59.
- Fuller, C. M., Biro, D. P., and Wilson, R. L. 2009. Decision support for determining veracity via linguistic-based cues. *Decision Support Systems* 46(3), 695-703.
- gCLUTO: Graphical Clustering Toolkit 1.2. Dept. of Computer Science, University of Minnesota.
- Hair, J.F., Anderson, R.E., Tatham, R.L. and W.C. Black. 1995. *Multivariate Data Analysis with Readings*. Upper Saddle River, NJ: Prentice Hall.
- Hancock, J. T., Curry, L. E., Goober, S., and Woodworth, M. 2008. On lying and being lied to: A linguistic analysis of deception in computer-mediated communication. *Discourse Processes*, 45(1), 1-23.
- Indyk, P. 1999. A Sublinear-time Approximation Scheme for Clustering in Metric Spaces. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science*.
- Karypis, G. 2003. Cluto: A Clustering Toolkit. Minneapolis: Univ. of Minnesota, Comp. Sci. Dept.
- Mann, W. C., and Thompson, S. A. 1988. Rhetorical Structure Theory: Toward a Functional Theory of Text Organization. *Text*, 8(3), 243-281.
- Manning, C.D. and H. Schütze. 1999. *Foundations of Statistical Natural Language Processing*. Cambridge, MA: MIT Press.
- Mihalcea, R., and Strapparava, C. 2009. The Lie Detector: Explorations in the Automatic Recognition of Deceptive Language. In *Proceedings of the ACL*, Aug. 2-7, Singapore.
- Pennebaker, J., and Francis, M. 1999. *Linguistic Inquiry and Word Count: LIWC*. Erlbaum Publisher
- Porter, S., and Yuille, J. C. 1996. The language of deceit: An investigation of the verbal clues to deception in the interrogation context. *Law and Human Behavior*, 20(4), 443-458.
- Rasmussen, M. and G. Karypis. 2004. gCLUTO: An Interactive Clustering, Visualization and Analysis System. *UMN-CS TR-04-021*.
- Rose, K. 1998. Deterministic Annealing for Clustering, Compression, Classification, Regression, and Related Optimization Problems. In *Proceedings of the IEEE* 86(11).
- Rubin, V.L. 2010. On Deception and Deception Detection: Content Analysis of Computer-Mediated Stated Beliefs. In *Proceedings of the American Soc. for Information Science and Tech. Annual Meeting*, Oct. 22-27, Pittsburgh.
- Rubin, V.L., and Conroy, N. 2011. Challenges in Automated Deception Detection in Computer-Mediated Communication. In *Proceedings of the American Soc. for Information Science and Tech. Annual Meeting*, Oct. 9-12, New Orleans.
- Rubin V.L., Conroy, N. 2012. Discerning Truth from Deception: Human Judgments and Automation Efforts. *First Monday* 17(3), <http://firstmonday.org>
- Salton, G. and M.J. McGill. 1983. *Introduction to Modern Information Retrieval*. New York: McGraw-Hill.
- Scholkopf, B. and A. Smola. 2001. *Learning With Kernels*. Cambridge, MA: MIT Press.
- Strehl, A., Ghosh, J. and R. Mooney. 2000. In *AAAI Workshop of Artificial Intelligence for Web Search*, July 30, 58-64.
- Taboada, M. 2004. *Building Coherence and Cohesion: Task-Oriented Dialogue in English and Spanish*. Amsterdam, Netherlands: Benjamins.
- Taboada, M. and W.C. Mann. (2006). Rhetorical structure theory: looking back and moving ahead. *Discourse Studies*, 8(3), 423-459.
- Tombros, A. 2002. *The effectiveness of query-based hierarchic clustering of documents for information retrieval*. PhD dissertation, Dept. of Computing Science, University of Glasgow.
- Vapnik, V. 1998. *Statistical Learning Theory*. NY: Wiley.
- Vrij, A. 2000. *Detecting Lies and Deceit*. NY: Wiley.
- Zhong, S. and Ghosh, J., 2004. A Comparative Study of Generative Models for Document Clustering. In *SIAM Int. Conf. Data Mining Workshop on Clustering High Dimensional Data and Its Applications*.
- Zhou, L., Burgoon, J. K., Nunamaker, J. F., and Twitchell, D. 2004. Automating Linguistics-Based Cues for Detecting Deception in Text-Based Asynchronous Computer-Mediated Communications. *Group Decision and Negotiation*, 13(1), 81-106.

Appendix A. Sample RST Analysis.




Appendix B1. Distributions of Expected Levels of Deception and Truthfulness in Deceptive Stories.

Legend: ▲ Expected level of Deception (Judges); ■ Expected Level of Truthfulness (Judges)
 — RST Level of Deception; - - - RST Level of Truthfulness (transformed to the interval (0,1) with 0 min)



Appendix B2. Distributions of Expected Levels of Deception and Truthfulness in Truthful Stories.

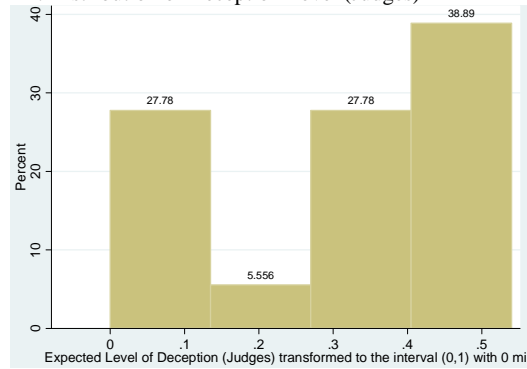


Appendix C. Comparison of the Normalized Frequencies of the RST Relationships in Truthful and Deceptive Stories: Difference in Means Test.

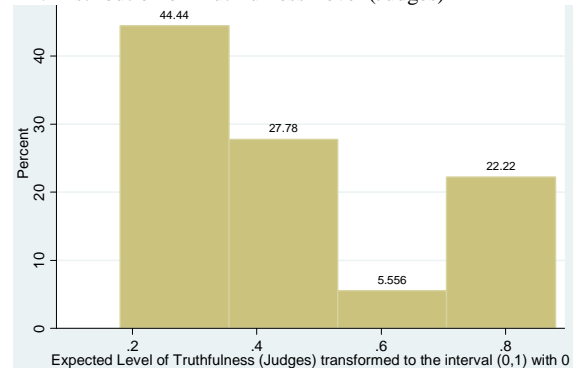
RST relationships appearing in truthful and deceptive stories with NO statistically significant differences	RST relationships appearing in the truthful stories with statistically significantly GREATER normalized frequencies than the deceptive ones	RST relationships appearing in the truthful stories with statistically significantly LOWER normalized frequencies than the deceptive ones
Background, Circumstance, Concession, Condition, Conjunction, Elaboration, Enablement, Interpretation, List, Means, Non-volitional cause, Non-volitional result, Purpose, Restatement, Sequence, Solutionhood, Summary, Unconditional	Antithesis (t=2.3299) Evidence (t=3.7996) Joint (t=1.5961) Volitional cause (t=1.8597) Volitional result (t=1.8960)	Preparation (t= -1.7533) Evaluation (t= -2.0762) Disjunction (t= -1.7850)

Appendices D1–D4. Distribution of Deception and Truthfulness Levels for Deceptive Stories

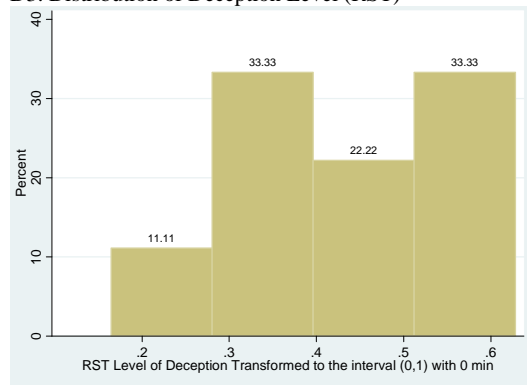
D1. Distribution of Deception Level (Judges)



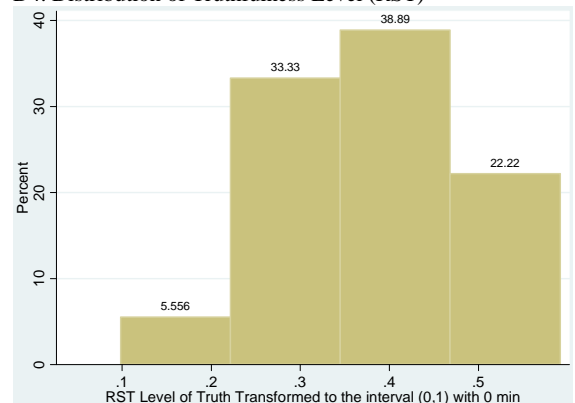
D2. Distribution of Truthfulness Level (Judges)



D3. Distribution of Deception Level (RST)

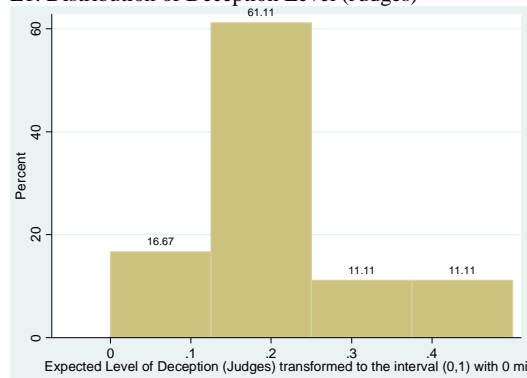


D4. Distribution of Truthfulness Level (RST)

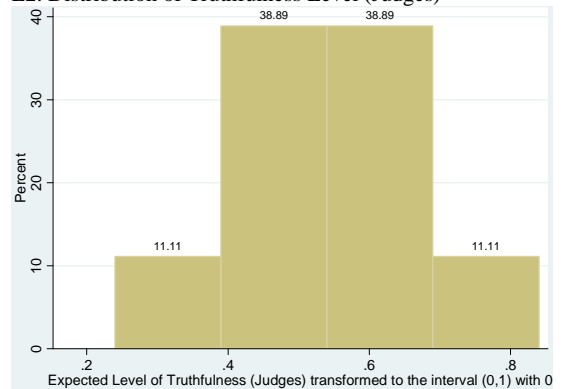


Appendices E1-E4. Distribution of Deception and Truthfulness Levels for True Stories

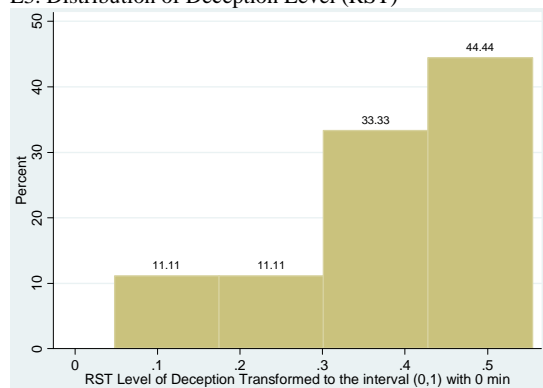
E1. Distribution of Deception Level (Judges)



E2. Distribution of Truthfulness Level (Judges)



E3. Distribution of Deception Level (RST)



E4. Distribution of Truthfulness Level (RST)

