A Cognitive Framework for Information Gathering with Deception Detection for Intelligence Analysis

Eugene Santos Jr., Qunhua Zhao, Gregory Johnson Jr., and Hien Nguyen
Computer Science and Engineering Department
University of Connecticut
371 Fairfield Road, U-155, Storrs, CT 06269-3155
{Eugene,qzhao,gjohnson,hien}@engr.uconn.edu

Paul Thompson
Department of Computer Science
Dartmouth College
6211 Sudikoff Laboratory, Hanover, NH 03755
Paul.Thompson@dartmouth.edu

Keywords: Search and Retrieval, Summarization, User Modeling, Information Assurance, Misinformation and Deception Detection

Abstract

Gathering relevant and credible information from massive, available data spaces is critical for intelligence analysis. We describe our ongoing effort in constructing a framework of services to support the analysts, which include personalized information retrieval and summarization with deception detection. By incorporating a user model that captures the analyst's intent, the system can retrieve information that is relevant to the analyst's information needs, and present the information properly as a summary according to the analyst's personal interests. By monitoring for a inconsistencies in the retrieved information, the system detects potential misinformation and/or disinformation. Our goal is to help the analysts effectively retrieve information based on their needs, and filter out or identify information that might be deceptive or misleading.

1. Introduction

Collecting information is one important step in the intelligence process. How to extract relevant information quickly from the massive amounts of available data and how to filter out unreliable or even misleading information are two major challenges in this step.

To meet the user's information need is the task of information retrieval (IR). Traditional IR applications rank the relevancy of information based on the matches between keywords found in queries and their frequency in the sources. However, individual difference has been ignored. Analysts may have different interests which vary according to their personal knowledge and experience even with the same task. In addition, they also have different cognitive search styles. Thus, individuals will apply varied relevance criteria in the searching process (Borlund 2003). Clearly, a user model of an intelligence analyst is the key to assisting his/her information retrieval (IR) task. User modeling (UM) has been employed to help improve users' IR performance since the early 80's (Brajnik et al. 1987). In our recent effort, a UM approach was developed to capture an analyst's intent in his/her IR process in order to better serve his/her information needs (Nguyen et al. 2004).

Document summarization is another effort for handling massive information (Mani et al. 1999). In an IR process, the analysts may heavily depend on the summary to decide if the original text is worth reading. Sometimes, they might even depend on the summaries to get the key ideas from the sources. Various approaches have been employed in developing automated summarization techniques. However, similar to the situation with IR applications, these systems ignore the user's unique information needs. Thus, it is possible for an intelligence analyst to overlook a document because its summary does not contain information relevant to his/her interests. Therefore, user-centered summarization becomes an increasingly important research area.

A summary may also need to be created from a set of documents instead of a single text. In this case, the breadth of the documents in the set makes it even more difficult to determine what should be covered in a good summary (McKeown et al. 2002). This can lead to a variety of problems such as incompleteness and omission, and cause critical problems in intelligence analysis. Automatically identifying document-sets that actually contain a diversity of information can help alert analysts not to take the summary at face value, and eventually leads to proper personalization of summaries with respect to user tasks and needs.

For intelligence, gathering relevant information is far from enough, because the accuracy and credibility of collected information are very important and can not be taken for granted. Detecting inconsistencies, either caused by intentional deception or through unintentional misinformation, appears to be useful in assisting intelligence analysts. With accurate and credible information, the analysts can be more certain about a particular hypothesis. Recognizing the inconsistency or even deceptive information can actually provide analysts valuable information to present to decision makers.

In this paper, we introduce our ongoing efforts in assisting intelligence analysts to battle the information overload and misinformation problems. We first briefly introduce our user modeling approach for IR, the use of document graphs for information extraction and representation; followed by our design of user model supported, personalized summarization. Next, we discuss how to help the analysts recognize possible deceptions or misinformation. Finally, we conclude with future work.

2. Intelligent Information Management to Overcome the Problem of Information Overload

Collecting relevant information is an important step in the intelligence process. Since information needs are individually different, UM is the key to meeting the challenge of determining relevance based on individualized criteria. Related to the efforts in dealing with information overload, document summarization tries to compress the information contained in a long text or a collection of related texts. Personalized summaries would enable an intelligence analyst to get the key ideas in the documents and relevant details, and help him/her to find the relevant documents more effectively.

2.1 User Modeling for Capturing Analyst Intent

User model is the key component in assisting intelligence analysts to meet the challenge of gathering relevant information from the massive amounts of data. The user model is situated between the intelligence analyst and the software tools, and is tailored to the specific analyst's needs to enhance and better support analytic activities (Nguyen et al. 2004; Santos et al. 2003). It continuously monitors the interaction between the analyst and the system tools, and proactively predicts and explains analyst's goals and intent, which serve as inputs to the various software tools. It also learns from a user's feedback to evolve to better model the analyst with respect to his/her tasks and specific domains of operation.

2.2 Intelligent IR Service Supported by IPC User Model

We developed a user model for IR services, which has three components: Interests, Preferences and Context (referred to as the IPC model) (Santos et al. 2003). Interests capture the individual's focus; Preferences capture how the input queries are modified and if the user is satisfied with the results; and, Context is what the user learned from the retrievals and provides insight into the user's knowledge. User Interests, Preferences and Context are captured in an Interest set, a Preference network and a Context network, respectively. Interest set is a list of concepts/entities, each of which is associated with an interest level. Each interest level is initially determined based on the current query, it is then updated based on the intersections of the retrieved relevant documents. The Preference network is represented as a Bayesian network (Pearl 1988), which consists of precondition nodes for the environment in which the user is searching for relevant information, goal nodes for tools that are used in query modification, and action nodes which represent how the user query should be modified. The Context network is a directed acylic graph that contains concept nodes and relation nodes. It is created dynamically by finding the intersection of retrieved relevant documents.

In an IR session, the user model tracks the analyst's intent, then modifies the query proactively in order to provide better retrieval results. Also, instead of searching by keywords, our IR system is based on document graphs (DGs). Each document in the collection is transformed into a DG, which is a directed graph of concepts/entities and the relations between them. It contains two node types, concept/entity nodes and relation nodes. Currently, only two relation types, "*isa*" and "*related to*", are captured for simplicity. DG construction is an automatic process, which includes tokenizing a text into sentences; parsing each sentence; extracting noun phrases from the parsing results; and generating relations based on heuristic rules (Santos et al. 2004).

A query is also transformed into a query graph (QG) by the same process. Next, the QG is matched against each DG in the collection. The similarity between QG to DG is given by the equation:

$$Similarity\ (QG, DG) = \frac{n}{2N} + \frac{m}{2M}$$

which is modified from (Montes-y-Gómez et al. 2000). N is the total number of concept/entity nodes in QG, and n is the number of matched concept/entity; M is the number of relation nodes in QG, and m is the number of matched relations. We say we find a matched relation only when both of the two concept/entity nodes linked to that relation node are matched, and the relation type is also matched. Since the size of a QG is usually significantly smaller that that of a DG, we used the number of concept/entity nodes and relation nodes in the QG as *N and M*, instead of the summation of node numbers in QG and DG.

Our IPC user model enhanced IR approach has been evaluated by comparing against Ide dec hi which is an approach using relevance feedback based on term frequency and inverted document frequency (TFIDF). Ide dec hi is still considered to be the best traditional IR approach. Our results showed that, based on the average precision at three point fixed recall (0.25, 0.5 and 0.75), our system based on DG significantly improved the retrieval performance, where average precision was 0.19 for TFIDF and 0.29 for our system (for original collection). UM helped further improve the retrieval performance and achieved compatible results to Ide dec hi (the average precisions were 0.31 for Ide dec hi and 0.33 for IPC) (Nguyen et al. 2004). As the Cranfield data has pre-classified relevancy on a 4-point numeric scale (1 for most relevant and 4 for least relevant) for all the relevant documents, we also found that our system retrieved significantly more highly ranked relevant documents than the TFIDF approach (Nguyen et al. 2004). We have also conducted a user evaluation study with actual end users, three working intelligence analysts, and compared our user model enhanced IR system with a commercial off-the-shelf system, the Verity Query Language. The experiment took place in a NIST laboratory, which positively concluded that our UM approach tracked the individual's interests, adapted to their individual searching strategies, and helped retrieve more relevant documents.

2.3 User-Centered Summarization

It is obvious that, from the same text, different people with different tasks will generate different summaries (Mani and Bloedorn, 1998). Deciding what should be selected from the original text is the key challenge in summarization. In our experiment with the Document Understanding Conference (DUC) 2002 collection for multi-document summarization, we identified two types of document-sets: sets which consist of closely related documents, and those of highly diverse texts. Intuitively, it is more difficult to create a good summary in the latter case. As DUC has classified document-sets into four categories, we found most of the document-sets in categories 1 and 2 belonging to type-1, and most of the documents in categories 3 and 4 belonging to type-2. Document-sets in categories 1 and 2 contain articles about a single disaster or event occurring within a seven day period, and are more similar to each other. Sets in category 3 are about distinct events with no time window restriction. In category 4, biographical information about a person is presented from different viewpoints. Therefore, topics in a document-set in categories 3 and 4 are more diverse than those in the other two categories. This difference could be detected by comparisons based on the average similarities of DGs for documents in sets. The average similarities among documents within a set are 0.10, 0.12, 0.06 and 0.07 for categories 1 to 4, respectively.

In the case of multi-document summarization, automatically identifying a summary which has been created from a document-set covering a breadth of topics can alert the analyst to possible bias or information omission. Also, the two types of document-sets actually reflect the requirements for summarization, that is, a common core of general topic with information from different aspects or details (branches) that are relevant to the core. Thus, a user-centered approach is critical to identifying what kind of "branching" information needs to be present in the summary to satisfy the specific user interests and ultimately address their individual tasks. In fact, related documents/reports collected for one intelligence analysis task usually contain different details or viewpoints. Therefore, selecting the proper pieces of information to represent in the summaries is important for helping the analysts identify the relevancy of the original texts, and ultimately connect the "dots".

As we discussed before, DG generation is a process of information extraction and representation. As a result, the important concepts, entities and relations in the text are captured in the DG. Currently, we are developing an automatic summarizer based on DG. First, we generate a DG core for a document(-set), which contains the most general topics covered in the content. For document sets, the core is constructed by majority vote; while for a single document, the core is created based on the weights of the relations. Next, the core is expanded by inserting relevant relations based on graph structure. The knowledge of the user's IR foci/interests learned by our IPC user model can be used to determine which relations relevant to the DG core would more likely satisfy each individual's information needs. An extract can then be generated based on this DG, where each sentence in the source text(s) are then compared to this DG based on the concepts/relations it contains, and those that cover more information will be selected.

3. Detecting Deceptive Information

Collected information must be processed before it can be used in intelligence analysis, which involves evaluating the accuracy and reliability of the information. Sound assessments can only be made based on reliable information, with the awareness of inconsistent data or potential deception. In this section, we describe our efforts to extend previous work in deception detection in intelligent systems to intelligence analysis as the basis for a tool for detecting deception in intelligence information.

3.1 Misinformation and Deception

Misleading information, misinformation, does not necessarily indicate deception, it can result from a multitude of sources including honest misinterpretation of the data, sensor noise, errors in communication and other sources of uncertainty in the intelligence gathering process. False information involved in a deception, disinformation, is intentionally communicated to mislead. Deception, as defined by Burgoon and Buller (1994), is a "deliberate act perpetrated by a sender to engender in a receiver beliefs contrary to what the sender believes is true to put the receiver at a disadvantage". Several key elements are included in this definition. First, deception is an intentional act; Second, deception is goal-directed; Third, influencing the decision making process and subsequent actions is achieved by providing information that is known to be incorrect. Intelligence analysts routinely encounter both misinformation and disinformation during their analysis task. In many cases, information is misleading at the least and deceptive in the worst case. Furthermore, an aggregate effect from a series of small deviations, which taken individually would still seem credible, would cause significant difference in what could be concluded (Cybenko et al. 2002). Tools which assist in identifying potential misinformation or disinformation can provide essential help to the analyst in information evaluation. Finally, the identification of deceptions is valuable for intelligence in itself. This knowledge may be useful in deception or counter-deception operations (Whaley 1969).

3.2 Detecting Deception

Our work follows the general model of deception detection proposed by Johnson and Grazioli (1993). The first step in this model is activation, which is the detection of something unexpected in the environment. The remaining steps serve to form hypotheses regarding suspected deceptions, correct knowledge of the environment and reason with the corrected knowledge.

Activation focuses on detecting discrepancies in the environment. It appears to be the most useful in assisting intelligence analysts. This approach makes no distinction between, and is therefore useful in detecting, potential misinformation as well as disinformation. Detecting discrepancies offers our best chance at detecting deceptions. Whaley and Busby (2002) note any deception leaves at least two clues. Discrepancies related to what is hidden and those related to what is shown in its stead. The authors note that every real thing has a large, but finite, number of characteristics and every imitation shares at least one and often many of these characteristics. However, no imitation is perfect. Even the most perfect clone lacks two characteristics – it is not the first and it has a different history (Whaley and Busby 2002). If either an added or deleted characteristic is detected, the manipulations are exposed. The most important corollary of this rule is that one need only discover one of the discrepancies, of either added or removed characteristics, to unravel the deception. In this case, the wealth of information available to contemporary intelligence analysts becomes an asset rather than a liability in information overload. The IR algorithms included in this framework complement deception detection in that they assist the analyst in considering all of the relevant, and only the relevant, information available.

The Ombudsman's Method (Whaley and Busby 2002) illustrates the value of this approach. Busby, a professional magician, was charged with a challenging task of teaching casino dealers how to detect cheating without teaching them how to cheat. Furthermore, Busby did not know the specific rules of the games. The successful method which resulted was named as the Ombudsman's Method. In 1984, Whaley, using the method, was able to determine how a magic card trick was performed, one which no professional magician aside from Busby was able to discover. They were also successful in teaching the method to a group of intelligence analysts. The analysts were then able to determine the method by which a complicated magic trick was done merely by viewing a video of the trick and being alert for discrepancies.

Our initial experiments in evaluating activation techniques focused on probabilistic intelligent systems. Both traditional expert systems and multi-agent systems were considered (Santos and Johnson 2004). In the single expert case, a model of the expert was constructed by observing expert responses, which was used to predict future responses. In the multi-expert case, we assume that the opinions of an intelligent agent are likely to be correlated with the opinions of other agents with similar knowledge. Thus, the availability of the opinions of other agents in the multi-agent system provides an opportunity to detect deception by comparing the agent opinions when they share knowledge over the same domain. Our results indicated that, when agent opinions are correlated, it is possible to predict expert opinions reasonably accurately for the purpose of detecting deception. Likewise, presuming correlation between documents on the same topic, we intend to compare the content of related documents to detect discrepancies. By analyzing the content of a set of documents for discrepancies in the information, we can detect misleading information which may be indicative of deception.

In the case of intelligence, we detect deception by analyzing the content of relevant documents to detect discrepancies. Further information can be supplied to the intelligence analyst by developing hypotheses regarding possible deceptions. Patterns in the identified discrepancies can be compared to the patterns of the six deception tactics proposed by Bell and Whaley (Bell 1991). In the case where a decision model exists, sensitivity analysis of the decision network can reveal the goal of the deception by identifying the alternate actions and goals that would have been taken as a result of the deceptive information. However, significant difficulties lie in the limited ability of the state of current information extraction technology to identify the nature of the relationships between concepts in a document. It may be necessary for user assistance to specify the type of relationships among concepts using templates or other techniques. From an information retrieval perspective, statistical information which tells how closely related the concepts can identify concepts central to the document set. From a user-centered point of view, we can determine if the deception is related to the focus of the user's search. This additional information is helpful in determining the potential impact of deceptive discrepancies in the retrieved information. The final step in the deception detection model proposed by Johnson and Grazioli involves continuing the prescribed task with corrected information. However, the presence of deception is also valuable information and should be included in the final analysis. Therefore, we augment the corrected or verified information with the knowledge that a deception is being perpetrated. Finally, collaborative deceit from several sources must be considered. In this case, it may be the case that the deceptive opinions will not only deviate from expectations, but will also be clustered in order to produce the desired effect. That is, all colluding information sources would tell the same lie in order to achieve the desired effect.

4. A Short Example

Intelligence analysis has been described as a process of "finding the dots" and "connecting the dots." (Hollywood et al. 2004). We envision the following scenario, which is based on (Hollywood et al. 2004): Imagine an analyst is required to investigate any correlation between reports of problems with tuna boats in Sydney, Australia and Singapore. The initial search for "tuna boats in Sydney and Singapore" produces the following summarizations:

1) Feb. 4 Twenty-four tuna boats are ordered in Seattle for export to Singapore under a Panamanian flag; export papers are filed.

2) Aug. 6 A Seattle paper announces completion of an order for 24 tuna boats to a Panamanian concern.

3) Aug. 9 Twenty-four Panamanian tuna boats arrive in Sydney, are registered in the harbor.

4) Sept. 8 An application is made for 18 tuna boat berths in Singapore's harbor by a Panamanian firm.

5) Oct. 4 CIA officer in Singapore reports tuna boats causing problems in the harbor.

6) Oct. 16 Maryland Coast Guard reports Panama-flagged boats acting unconventionally and being uncooperative in Baltimore, and possible related events at Philadelphia.

7) Oct. 21 Commerce Section of U.S. Embassy in Singapore reports notable commercial activity by a tuna boat company.

8) Oct. 23 Economics officer at U.S. Embassy in Singapore notes non-routine hiring practices at the harbor, new tuna boat activity, and potential for economic recovery.

9) Oct. 24 Miami DEA reports possible relationship between local intercepted tuna boats and activity in Baltimore and Philadelphia.

10) Nov. 4 CIA in Singapore reports coincidence of problems with tuna boats in Sydney and Singapore.

The analyst indicates to the IR system, that these reports are relevant to his task. There are discrepancies in the information in reports 1) and 2) when compared with that in reports 3) and 4). Deception detection algorithms detect the discrepancy between the quantity of 24 boats ordered and delivered from a Seattle firm and the 18 boats accounted for in Singapore's harbor. Six boats are unaccounted for, this is categorized as a Masking Dissimulative tactic (Bell 1991). The location of six of the boats is concealed. There is also a discrepancy in that the boats were ordered for export to Singapore, however, the boats were delivered to Sydney, Australia. This is classified as a Decoying Simulative tactic. The export papers were filed with deceptive information.

These reports originated from several different agencies (e.g. CIA, FBI) and the IR tools were able to "find the dots" among the multitude of intelligence reports generated by various agencies. The context network created in the UM process, which represents the relationships between concepts, provides the connections between the dots. Since the user model learns from user feedback (the documents the analyst marked as relevant) that these tuna boats are related to a Panamanian company, a modified query that includes this information retrieves two additional reports,

11) Oct. 1 A Panamanian firm requests off-loading privileges at Baltimore public wharf.

12) Oct. 10 A wharf in Philadelphia is leased to a Panamanian firm.

Subsequent modified queries retrieve additional relevant documents and additional dots are found and connected. The analyst decides that there is sufficient evidence to warrant further investigation of the cargo activity of tuna boats in Baltimore and Philadelphia. From the example, we can see that better IR results can help the analysts make links among the evidence. Also, a good cluster of documents/reports can make it easier to identify inconsistent information. The results can then lead to the verification of obtained information and guide further collection and analysis.

5. Conclusions and Future Work

Our goal is to develop a framework that helps the intelligence analysts to retrieve information that is relevant to their individual interests as well as their current task while detecting possible misinformation and deception. We view personalization as the solution, and user modeling as the key element, for the IR process. The IPC user model that has been briefly described above has been successfully deployed in an IR application. Several evaluations have demonstrated that the IPC user model was able to follow the user's interests, and improve the IR performance. Also, by incorporating a user model into the process of automatic document summarization, it will enable us to provide summaries that are geared to the personal information needs of the individual analyst. Such summaries can assist analysts to identify relevant documents/reports more effectively.

Although user modeling enhanced IR and summarization can provide a way to reduce the problem of information overload in selecting relevant information, the analysts still face the problem of judging the reliability of the information they get. Detecting deceptive information thus becomes vital to making a correct decision, or at least a well informed one. This applies to unintentional misinformation as well as intentional disinformation.

Following the model of deception detection of Johnson and Grazioli (1993), the first step is to monitor for inconsistencies among the attributes in the environment, which we consider to be the most useful approach in detecting misinformation and disinformation. Significant work is needed to improve the methods for identifying the relationships between concepts in the captured information related to time, places, names, etc., where the advance in Automatic Content Extraction (ACE) program (http://www.nist.gov/speech/tests/ace/) seems promising. Public sources can be used to extract relevant information for testing the detection of inconsistence. Continuing work will also include determining the minimum number of information sources needed to accurately detect deception, and empirical studies to reveal the proper threshold that can be used to identify when an information source is not correlated enough with expectations. Further studies will focus on incorporating a decision model to track possible deviates in the conclusions. Based on evaluation of the explanation of why such deviation happens, we can identify possible deception and even the deceiver's goals.

Acknowledgments

This work was supported in part by AFOSR Grant No. F49620-03-1-0014 and NGA NURI Grant no HM1582-04-1-2027).

References

Bell, J. Bowyer and Whaley, B. 1991. *Cheating and Deception.* Transaction Publishers, New Brunswick, NJ.

Borlund, P. 2003. The concept of relevance in IR. *Journal of the American Society for Information Science and Technology* 54 (10): 913-925.

Brajnik, G., Guida, G. and Tasso, G. 1987. User Modeling in Intelligent Information Retrieval. *Information Processing and Management* 23(40): 305-320.

Burgoon, J. K. and Buller, D. B. 1994. Interpersonal Deception: III. Effects of Deceit on Perceived Communication and Nonverbal Behavior Dynamics, *Journal of Nonverbal Behavior* 18(2): 155-184.

Cybenko, G., Giani, A. and Thompson, P. 2002. Cognitive Hacking: A Battle for the Mind. *Computer*, August 2002, 50-56.

Johnson, P. and Grazioli, S. 1993. Fraud Detection: Intentionality and Deception in Cognition. Accounting, *Organizations and Society* 25: 355–392.

Hollywood, J., Snyder, D., McKay, K. and Boon, J. 2004. Out of the Ordinary: Finding Hidden Threats by Analyzing Unusual Behavior. RAND Corporation, Santa Monica, CA.

Mani, I. and Bloedorn, E. 1998. Machine Learning of Generic and User-Focused Summarization. In *Proceedings of the 15th National Conference on AI (AAAI-98)*. Madison, WI, USA. 26-30 July, 1998, 821-826.

Mani, I., House, D., Klein, G., Hirschamn, L., Firmin, T. and Sundheim, B. 1999. The TIPSTER SUMMAC Text Summarization Evaluation. In *Proceedings of the 9th Conference on European Chapter of the Association for Computational Linguistics.* Bergen, Norway. 8-12 June, 1999, 77-85.

McKeown, K., Barzilay, R., Blair-Goldensohn, S., Evans, D., Hatzivassiloglou, V., Klavans, J., Nenkova, A., Schiffman, B. and Sigelman, S. 2002. The Columbia Multi-Document Summarizer for DUC 2002. *ACL-2002: Workshop on Automatic Summarization (including DUC 2002)*. Philadelphia, PA, USA. 11-12 July, 2002.

Montes-y-Gómez, M., Gelbukh, A. and Lópes-López, A. 2000. Comparison of Conceptual Graphs. *The 1st Mexican International Conference on Artificial Intelligence*. Acapulco, Mexico, April, 2000, 548-556.

Nguyen, H., Santos, E. Jr., Zhao, Q and Wang, H. 2004. Capturing User Intent for Information Retrieval. In *Proceedings of the 48th Annual Meeting for the Human Factors and Ergonomics Society*, New Orleans, LA, USA. 20-24 September, 2004. 371-375.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA

Santos, E. Jr. and Johnson, G. 2004. Toward Detecting Deception in Intelligent Systems. In *Proceedings of the SPIE: Defense & Security Symposium, Vol. 5423.* Orlando, FL, 2004. 131-140.

Santos, E. Jr., Mohamed, A.A. and Zhao, Q. 2004. Automatic Evaluation of Summaries Using Document Graphs. In Proceeding of ACL 2004, Workshop on Text Summarization Branches out, Barcelona, Spain. 21-26 July, 2004, 66-73

Santos, E. Jr., Nguyen, H., Zhao, Q., and Wang, H. 2003. User Modelling for Intent Prediction in Information Analysis. In *Proceedings of the 47th Annual Meeting for the Human Factors and Ergonomics Society*. Denver, CO, 2003. 1034-1038.

Whaley, B. 1969 *Strategem: Deception and Surprise in War*. MIT Press, Cambridge, MA.

Whaley, B. and Busby, J. 2002. Detecting Deception—Practice, Practitioners, and Theory. Godson, R. and Wirtz, J. J. (eds.) *Strategic Denial and Deception: The Twenty-First Century Challenge*, Transaction Publishers.