International Journal of Latest Research In Engineering and Computing (IJLREC)
Volume 1, Issue 1 : Page No.33-39, September-October 2013
www.ijlrec.com

An Analysis of Fuzzy Approach For Detecting Anomalous Behaviour with E-mail Traffic

Abhishek Shukla[1], Kranti Deo Singh[2]

[1]Research Scholar Singhania University IT Department
R.D.Engineering College, Duhai,Ghaziabad (India)
[2]Asstt. Prof. HRIHM,Ghaziabad

Abstract: This paper investigates the use of fuzzy inference for detection of abnormal changes in email traffic communication behaviour. Several communication behaviour measures and metrics are defined for extracting information on the traffic communication behaviour of email users. The information from these behaviour measures is then combined using a hierarchy of fuzzy inference systems, to provide an abnormality rating for overall changes in communication behaviour of suspect email accounts. The use of fuzzy inference is then demonstrated with a case study investigating the email traffic behaviour of a person's email accounts from the Enron email corpus.

Keywords: **Email, traffic analysis, electronic surveillance, anomaly detection, fuzzy logic, fuzzy inference system, abnormality ranking, Enron email corpus.**

INTRODUCTION

On 10th August 2006, 21 terror suspects were arrested in Britain on suspicion of plotting to blow up United States bound commercial air flights with liquid explosives (Natta et al., 2006). It was reported that British security services, MI5, had been monitoring these suspects for up to at least 12 months prior to making the arrests in August 2006. The New York Times (Natta et al., 2006) reported that MI5 had used several sources of information to monitor the activities of the British terror suspects. These methods included: bugging their apartments, tapping their phones, monitoring their bank transactions, and eavesdropping on their Internet traffic and email messages.

This British terror case highlights the importance of monitoring the activities of terror suspects. Monitoring helps law enforcement investigators keep track of what terror suspects are doing, as well as who they are communicating with, and whether suspects are doing anything that indicates an unusual change in their pattern of behaviour compared to their normal activities (e.g. informing terror cell members when to conduct the attack). If the British security services had not been keeping watch on the activities of the British terror suspects and made the arrests based on what they had observed, the world might have experienced another airline related tragic event, similar to the terrorist attacks in the United States on September 11, 2001 (Whitney and Strasser, 2004).

Another point to note from the New York Times article is how the use of multiple sources of information by British Security Services may have helped to provide a broader perspective on what the terror suspects were doing. Multiple sources of information such as phone tapping, monitoring of bank transactions, and eavesdropping on Internet traffic and email messages, provided the British security services with a variety of sources for detecting any unusual patterns of behaviour or change from normal habits (e.g. an unusually large bank withdrawal). One of the difficulties in dealing with multiple sources of information is how to combine or "fuse" the information together. Some of the information sources may show evidence that unusual activity is occurring, but sometimes it may not be clear to the investigator how to combine the information together. Another problem is that it may be difficult for the investigator to know which monitored suspect should be observed more closely either as a matter of priority or based on the available evidence.

Our team is to analysis of email traffic communications, with a focus on determining how artificial intelligence techniques could be useful in aiding the user/intelligence analyst to investigate a suspected individual's email traffic communication behaviour. In our previous work (Lim et al., 2005, Lim et al., 2006) an email traffic analyser system was developed as a conceptual system to investigate the use of data visualisation techniques and decision trees (Witten and Frank, 2005, Negnevitsky, 2004) for finding "unusual" communication behaviour from simulated email traffic data. Our team focuses on developing a new anomaly detection module for the email traffic analyser system, which analysis a list of suspects for deviations from their normal patterns of communication behaviour in email traffic and alerts the user when an abnormal change

in communication behaviour has occurred. The recent work also looks at what the email traffic analyser system can reveal from genuine email traffic data. A diagram of the email traffic analyser system is shown in Figure 1.

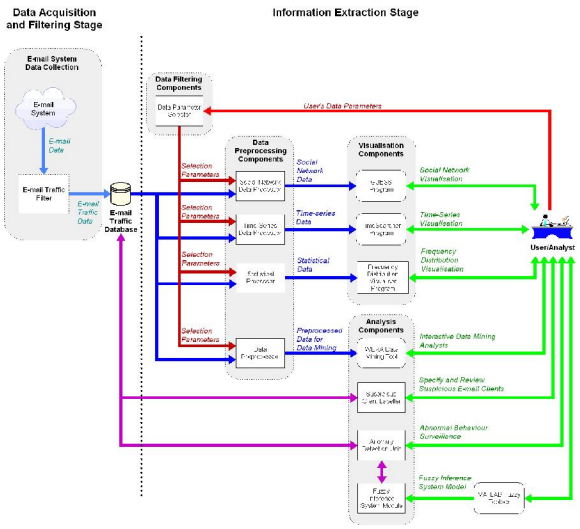

Figure 1: The email traffic analyser system.

In this paper, a brief description is first provided on anomaly detection and how the method of anomaly detection is being used to detect changes in email traffic communication behaviour. The second part of the paper
describes defining email traffic communication behaviour measures and how these will be used to record behavioural information on the email user being analysed. The third part of the paper describes how the anomaly detection module will profile the behaviour of email users and detect changes in communication behaviour patterns. The fourth section describes how fuzzy inference is being used to combine information from
different communication behaviour measures. This is then followed by a case study of the Enron email corpus, comparing the alert results produced by individual communication behaviour measures and the results produced
after fusing the information together using fuzzy inference.

ANOMALY DETECTION

The main aim of our current work is to monitor the email traffic of a suspected individual for any significant deviations from their normal communication behaviour patterns. The purpose of this is to bring to the attention of the user/analyst that an abnormal or unusual event is occurring and assist them in finding the location of the unusual event in the data. Our aim is to just inform the user about the presence of an unusual change in communication behaviour for the monitored suspect and allow the user to utilise data visualisation tools (Lim et al., 2005, Lim et al., 2006) or other analysis tools to investigate the details of that unusual event. We leave it up to the user to decide the context or meaning of the unusual event (e.g. is it a planned terrorist attack or a planned birthday party?), rather than try to encode the contextual knowledge into the system.

The method being used to detect changes in email traffic communication behaviour is anomaly detection, a method that is commonly used in intrusion detection (Bace and Mell, 2001) to detect new types of intrusion attacks, previously unknown to a computer system or computer network. Anomaly detection is based on the idea
that the computer system or computer network has a "normal" operating state, which can be used to determine if the system is currently under attack from an unknown intruder. In intrusion detection, the intrusion detection system (IDS) builds a model of the target computer system's "normal" state of behaviour and uses that model to
determine if the current state of the system is exhibiting significant deviations from the normal state of behaviour. If there are significant deviations, then the IDS informs the system or network administrator that there is an abnormal change in behaviour, indicating a possible attack on the computer system or computer network.

Although anomaly detection is commonly used in computer network security (Mohay, 2003), the same principles may also be applied for electronic surveillance applications when monitoring suspected individuals for changes in communication behaviour. In our email traffic analyser system, the anomaly detection module is used to detect possible changes in email traffic communication behaviour for a list of suspected individuals. The email traffic analyser system firstly requires the user to select a list of suspect email addresses from the email System being analysed, to specify which email accounts will be monitored. The user then selects a historical period of time or "profiling period" (e.g. a period of 1 year, starting at two years ago), which is used by the anomaly detection module to build behaviour profiles for all suspects and record their "normal" communication

behaviour patterns. After normal behaviour profiles have been created and stored in the email traffic database (Figure 1), the user then selects a recent period of time or "surveillance period" (e.g. a period of 6 months,ending on last week), which is used by the anomaly detection module to determine whether the recent behaviour of the suspects has significantly deviated from their "normal" communication behaviour.

DEFINING E-MAIL TRAFFIC COMMUNICATION BEHAVIOUR MEASURES

Before changes in communication behaviour patterns can be detected, communication behaviour measures need to be defined in order for the anomaly detection module to determine what kind of information will be used to record a change in communication behaviour. Thus, it is necessary to define communication behaviour measures, in order to describe particular aspects of an individual's email traffic communication behaviour and to describe how that individual's communication behaviour may have changed at
different periods of time. In this work, communication behaviour measures can be defined based on three sets of information taken from the header segments of email messages: the sender (the "from" field), the recipient/s (the "To", "CC", and "BCC" fields), and the date/time that the message was sent (from the "date" field). Using these three basic sets of information from the header component of email messages (excluding the content of email messages), the following types of communication behaviour measures can be defined:

A. E-mail Traffic Volume –

It is based on a count of the number of e-mails generated by an individual per hour, per day, per week, or per month, and sent to a particular contact. This provides information on the traffic volume flow of e-mails generated by an individual and the rate at which messages are being sent to particular contacts.

B. Delays Between E-mails Sent (or ªSending Delaysº) –

It is based on a measure of the time delays between each e-mail message sent by an individual. This provides information on expected delays between each message sent by an individual to particular contacts.

C. Replying Response Time (or ªReplying Delaysº) –

It is based on a measure of the time it takes for an individual to write a response e-mail to messages received from particular associates. This provides information on how quickly an individual is expected to reply to particular associates.

After defining the above communication behaviour measures, a set of metrics can be computed to produce a number that describes and summarises information about a particular communication behaviour measure. Each metric computed will provide information about an aspect of the monitored individual's communication behaviour. The following set of metrics was defined to describe and summarise each of the above communication behaviour measures, using statistical methods (Salkind, 2004, Gravetter and Wallnau, 2004, Chatfield, 1996):

D. Consistency of Weekly E-mail Traffic Volume -

It computes the autocorrelation of the weekly volume of e-mails produced by an individual, to indicate how ªconsistentº or ªreliableº an individual is with the weekly volume of e-mail traffic sent to particular associates. The autocorrelation, produces a number between ±1.0 to +1.0 to indicate the relationship between each time-series point in the weekly e-mail traffic volume data.

E. Percentage of Weekly E-mail Traffic Volume –

It computes the average percentage of e-mails sent to particular associates each week (e.g. 10% of e-mails per week to contact A, 40% per week to contact B, 50% per week to contact C).

F. Median of Sending Delays –

It computes the most commonly occurring time delays between e-mails sent to a particular associate, by using the statistical median.

G. Median of Replying Delays –

It computes the most commonly occurring response delay between e-mails replied to a particular associate, by using the statistical median.

It should be noted that when analysing e-mail traffic, one could also analyse the flow of e-mail messages in terms of the direction of the e-mail traffic (i.e. e-mail messages are either being sent or received by an individual). By taking the direction of e-mail traffic into account, the original four sets of metrics described above can be expanded into nine metrics, which summarises and describes an individual's incoming or outgoing e-mail traffic communication behaviour with each of their contacts.

The diagram in Figure 2 shows the mapping of the nine metrics in relation to the communication behaviour measures. Note that the metric titled ªMedian Of Combined Replying Delays With Contacts" considers the most commonly occurring response delay for both incoming and outgoing email traffic, hence providing information about the speed of the send response interactions between the individual and a particular associate.

These nine metrics are being used to record information about the state of the suspected individual's traffic communication behaviour patterns for the anomaly detection module. Note that the above is not an exhaustive list of all possible communication behaviour measures or metrics that can be extracted from email header information (i.e. sender, recipient, date/time information). The list defined above is the basic set of email traffic behaviour measures that we have chosen to focus upon for this work.

Other researchers working on similar or related email surveillance applications have explored different types of measures that can be extracted from sender, recipient, and date/time information. In the work by Stolfo et al. (2003a, 2003b), they have taken a pattern based or habit based approach where they consider particular habits of
email users, such as defining a measure for the time of day the user normally sends emails and a measure for the frequency of communication with particular contacts ("recipient frequency"). Another approach considered are ratio based measures, where Jiang et al. (2005) defined measures such as: ratio of new addresses vs. former addresses (measuring the rate that new email addresses appear), ratio of new senders vs. former senders (measuring the rate that new sending addresses appear), ratio of emails sent over time (measuring the volume of emails sent). Additional email traffic behavioural measures can be defined by using other header information fields (Tanenbaum, 2003) such as text/HTML formatting of the email, presence of attachments, or MIME file attachment type (Martin et al., 2005).

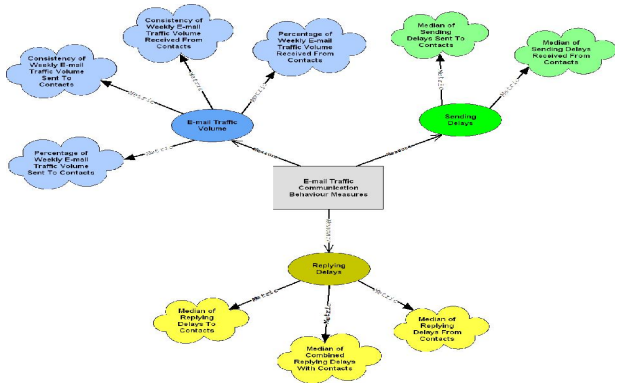

Figure 2: Mapping of the different patterns of behaviour that we are measuring from email message headers.

ANALYSING FOR CHANGES IN COMMUNICATION BEHAVIOUR

After the nine metrics were defined, these were used to build "normal" behaviour profiles for each of the suspect Email accounts during their profiling period. To build the normal behaviour profiles, each of the suspect's communication links with an associate is analysed and the nine metrics are computed for each communication link, which are then stored as the suspect's behaviour profile in the email traffic database. Figure 3 shows how the nine metrics are computed for each communication link with particular associates.

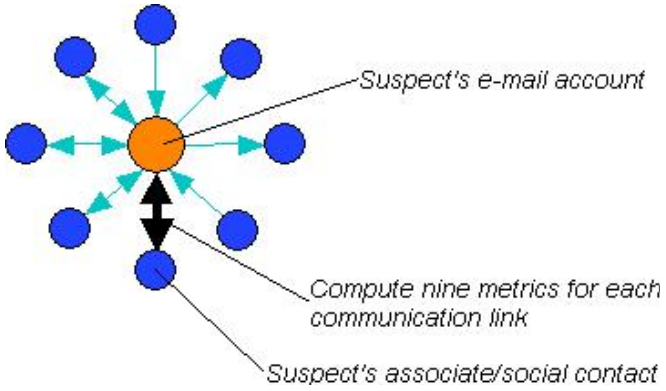

Figure 3: Diagram of how nine metrics are computed for each of the suspect's communication links.

To detect a change in communication behaviour, the nine
metrics are again computed for each of the suspect's communication links during the surveillance period and the recent communication behaviour measurements are compared with the measurements from the profiling period. If the recent behaviour of any communication link shows significant deviations from their previous communication behaviour patterns, then the user is alerted to the presence of an abnormal change in behaviour. In addition to alerting the user about changes in communication behaviour, the anomaly detection module also informs the user if there are new associates that have appeared in the surveillance period, which were not present in the suspect's "normal" behaviour profiling period.

The work by (Jiang et al., 2005, Stolfo et al., 2003a, Stolfo et al., 2003b, Martin et al., 2005) focuses on providing information on deviations in behaviour for each of the communication behaviour measures that they record from email users. However, the problem with their work is that they present the user/administrator/analyst with a lot of information about each of their communication behaviour measures, but do not summarise the email accounts that exhibit the most deviation in communication behaviour. For the user,all of the communication behaviour measures presented may be quite useful, but on first glance there is too much information for them to determine which email account is exhibiting the most deviation in communication behaviour and maybe thus the most interesting. Summarising all of the suspect email accounts' change inbehaviour is important, because if the user is trying to analyse the data for a large number of email accounts (e.g. more than 10), which email account should they pay attention to first? Which communication links should receive first priority in the investigation?

COMBINING INFORMATION USING FUZZY INFERENCE TECHNIQUES

To summarise the changes in communication behaviour of suspect email accounts, we investigate the use of fuzzy inference techniques. Fuzzy inference is a technique that employs the use of a concept called fuzzy logic (Zadeh, 1965). This is an artificial intelligence technique used to assist the computer to interpret vague or uncertain terms. As humans, we often use vague terms to describe things that we observe in the world around us, e.g. "the weather is hot", "that man is tall", "the danger risk is high". Computers normally cannot understand vague terms and must compute observations using crisp numbers, e.g. "the weather is 37.5ºC", "that man is 182 cm tall", "the danger risk is 89%". Fuzzy logic helps computers to interpret vague or uncertain terms in a similar manner to the way humans do, through the use of fuzzy sets (Zadeh, 1965).

Fuzzy inference builds upon the use of fuzzy logic and fuzzy sets (Mamdani and Assilian, 1975, Negnevitsky,2004), using fuzzy heuristic rules that encode knowledge using vague or uncertain terms. For example: "IF temperature is hot, THEN air conditioner output is high", "IF temperature is warm, THEN air conditioner output is medium". Fuzzy inference systems operate by processing input data that is crisp (e.g. 37.5ºC), interpreting that value by "fuzzifying" it (e.g. 37.5ºC is a member of the term "hot"), applying the fuzzy rules to determine the output (e.g. air conditioner output is high), then "defuzzifying" the output to produce a crisp number (e.g. air condition output level = 90%). One of the advantages of fuzzy inference is that it is able to process data that contains uncertain information and also has the ability to process input from several measurement sensors in parallel. Fuzzy inference is often used in decision support systems (Turban and Aronson, 2001) to provide advice on things that contain a level of uncertainty or risk, such as, for example, real estate evaluation (Bagnoli and Smith, 1998).

For the anomaly detection module, we use a hierarchy of several fuzzy inference systems, shown in Figure 6, to combine the input measurements from the nine communication behaviour metrics, and output a e commendation for the overall deviation in communication behaviour for each communication link (i.e. between the suspect and an associate). The final output recommendation given by the fuzzy inference hierarchy produces a number in the range of 0.0 to 1.0, where numbers close to 0.0 signify very little change in overall communication behaviour and numbers close to 1.0 signify a very large change in overall communication behaviour.

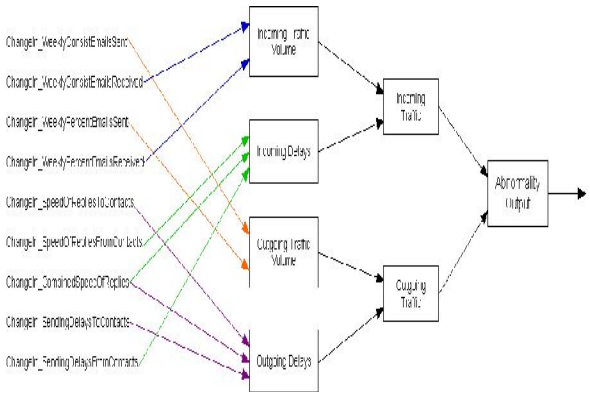

Figure 6: The fuzzy inference hierarchy used for the anomaly detection module, where each block is a fuzzy inference system.

CONCLUSION

We have shown how using fuzzy inference techniques may make the email traffic anomaly detection results easier for the user/analyst to interpret, through ranking the degree of abnormality for different communication links between the suspect and their associates. Most APPROACHES SHOWN BY other researchers, focus on presenting the user a whole selection of information on different communication behaviour measures, but do not provide a ranking for the user/analyst to decide which email addresses or communication links receives higher priority in the investigation of anomalous behaviour. The advantage of fusing together information from different communication behaviour measures to perform email traffic anomaly detection, and investigating a person's traffic communication behaviour from the Enron email corpus was also shown. Future work will involve comparing the results from the analysis of our simulated email data and real email data, investigating the use of different input grouping combinations for the fuzzy inference hierarchy, and investigating different time durations for the profiling and surveillance of the email user's traffic behaviour.

REFERENCES

1. Adar, E. (2006), GUESS: A Language and Interface for Graph Exploration. CHI 2006: The SIGCHI Conference on Human Factors in Computing Systems, Montréal, Québec, Canada, April 22 27:791-800.
2. Aris, A., Khella, A., Buono, P., Shneiderman, B. and Plaisant, C. (2005), Time Searcher 2, Human Computer Interaction Laboratory, Computer Science Department, University of Maryland, URLhttp://www.cs.umd.edu/hcil/timesearcher/, Accessed: 31st July, 2006.
3. Bace, R. and Mell, P. (2001), NIST Special Publication 80031:Intrusion Detection Systems, National Institute of Standards and Technology (NIST), URL http://csrc.nist.gov/publications/nistpubs/80031/sp80031.pdf, Accessed: 26 February 2004.
4. Bagnoli, C. and Smith, H.C. (1998) The Theory of Fuzzy Logic and Its Application to Real Estate Valuation.Journal of Real Estate Research, 16(2).
5. Bezdek, J.C. (1981) Pattern Recognition with Fuzzy objective Function Algorithms, Plenum Press, New York.
6. Chatfield, C. (1996) The Analysis of Time Series: An Introduction, 5th edn, Chapman and Hall, London.
7. Cohen, W.W. (2004), The CMU Enron Email Dataset, URL http://www.cs.cmu.edu/~enron/, Accessed: 12 October 2006.
8. Dickerson, J.E., Juslin, J., Koukousoula, O. and Dickerson, J.A. (2001) Fuzzy intrusion detection. Proceedings Joint 9th IFSA World Congress and 20th NAFIPS International Conference, 3: 1506-1510.
9. Diesner, J., Frantz, T.L. and Carley, K.M. (2005) Communication Networks from the Enron Email Corpus "It's Always About the People. Enron is no Different". Computational & Mathematical Organization Theory,11(3): 201- 228.
10. Fiore, A. and Heer, J. (2005), UC Berkeley Enron Email Analysis, URL http://bailando.sims.berkeley.edu/enron_email.html, Accessed: 12 October 2006.
11. Fox, L. (2003) Enron: The Rise and Fall, John Wiley & Sons, Hoboken, NJ.
12. Fusaro, P.C. and Miller, R.M. (2002) What Went Wrong at Enron: Everyone's Guide to the Largest Bankruptcy in U.S. History, John Wiley & Sons, Hoboken, NJ.
13. Gravetter, F.J. and Wallnau, L.B. (2004) Statistics for the Behavioral Sciences, 6th edn, Thomson/Wadsworth, Australia.
14. Jiang, T., Kim, W., Lhee, K. and Hong, M. (2005), Email worm detection using the analysis of behavior, in Distributed Computing and Internet Technology, Proceedings, SpringerVerlag Berlin, Berlin, 3816: 348-356.
15. Lim, M.J., Negnevitsky, M. and Hartnett, J. (2005), Tracking and Monitoring Email Traffic Activities of Criminal and Terrorist Organisations Using Visualisation Tools. 6th Australian Information Warfare & Security Conference, Geelong, Victoria, Australia, 24th to 25th November 2005: 112 -124.
16. Lim, M.J.H., Negnevitsky, M. and Hartnett, J. (2006) Personality Trait Based Simulation Model of the Email System. International Journal of Network Security, 3(2): 164-182.
17. Mamdani, E.H. and Assilian, S. (1975) An experiment in linguistic synthesis with a fuzzy logic controller. International Journal of Man Machine Studies, 7(1): 1-13.
18. Martin, S., Sewani, A., Nelson, B., Chen, K. and Joseph, A.D. (2005), Analyzing Behaviorial Features for Email Classification. Second Conference on Email and Anti Spam (CEAS 2005), Stanford University, Palo Alto, CA, July 21-22.
19. Mohay, G.M. (2003) Computer and Intrusion Forensics, Artech House, Boston.
20. Natta, D.V., Sciolino, E. and Grey, S. (2006), Details Emerge in British Terror Case, New York Times, URL http://www.nytimes.com/2006/08/28/world/europe/28plot.html?ex=1160712000&en=2582c7f1e971edd7 &ei=5070, Accessed: 12 October 2006.
21. Negnevitsky, M. (2004) Artificial Intelligence: A Guide to Intelligent Systems, 2nd edn, Addison Wesley, Essex.
22. Salkind, N.J. (2004) Statistics for people who (think they) hate statistics, 2nd edn, Sage Publications, Thousand Oaks, CA.
23. Shetty, J. and Adibi, J. (2005), The ISI Enron Database Schema and Brief Statistical Report, URL http://www.isi.edu/~adibi/Enron/Enron.htm, Accessed: 12 October 2006.

24. Stolfo, S.J., Hershkop, S., Wang, K., Nimeskern, O. and Hu, C.W. (2003a), Behavior profiling of email, in Intelligence and Security Informatics, Proceedings, SpringerVerlag Berlin, Berlin, 2665: 74-90.

25. Stolfo, S.J., Hershkop, S., Wang, K., Nimeskern, O. and Hu, C.W. (2003b), A behaviour based approach to securing email systems, in Computer Network Security, SpringerVerlag Berlin, Berlin, 2776: 57-81.

26. Tanenbaum, A.S. (2003) Computer Networks, 4th edn, Prentice Hall PTR, Upper Saddle River, NJ.

27. Turban, E. and Aronson, J.E. (2001) Decision Support Systems and Intelligent Systems, 6th edn, Prentice Hall,Upper Saddle River, NJ.

28. Whitney, C.R. and Strasser, S. (2004) The 9/11 Investigations : Staff Reports of the 9/11 Commission : Excerpts from the House Senate Joint Inquiry Report on 9/11 : Testimony from Fourteen Key Witnesses, Including Richard Clarke, George Tenet, and Condoleezza Rice, 1st edn, PublicAffairs, New York.

29. Witten, I.H. and Frank, E. (2005) Data Mining: Practical Machine Learning Tools and Techniques, 2nd edn, Morgan Kaufmann, San Francisco, Calif.

30. Zadeh, L. (1965) Fuzzy Sets. Information and Control, 8(3): 338-353.

