Chapter 7
Security Data Mining: A Survey Introducing Tamper-Resistance

Clifton Phua and Mafruz Ashrafi

Abstract Security data mining, a form of countermeasure, is the use of large-scale data analytics to dynamically detect a small number of adversaries who are constantly changing. It encompasses data- and results-related safeguards; and is relevant across multiple domains such as financial, insurance, and health. With reference to security data mining, there are specific and general problems, but the key solution and contribution of this chapter is still tamper-resistance. Tamper-resistance addresses most kinds of adversaries and makes it more difficult for an adversary to manipulate or circumvent security data mining; and consists of reliable data, anomaly detection algorithms, and privacy and confidentiality preserving results. In this way, organisations applying security data mining can better achieve accuracy for organisations, privacy for individuals in the data, and confidentiality between organisations which share the results.

7.1 Introduction

There is the exceptional progress in networking, storage and processor technology; as well as the increase in data sharing between organisations. As a result, there is the explosive growth in the volume of digital data, a significant portion of which is collected by an organisation for security purposes.

This necessitates the use of security data mining to analyze digital data to discover actionable knowledge. By actionable, we mean that this new knowledge improves the organisation's key performance indicators, enables better decision-making for the organisation's managers, and provides measurable and tangible results. Instead of purely theoretical data-driven data mining, more practical domain-driven data mining is required to discover actionable knowledge.

Clifton Phua, Mafruz Ashrafi
A*STAR, Institute of Infocomm Research, Room 04-21 (+6568748406), 21, Heng Mui Keng Terrace, Singapore 119613, e-mail: {cwphua,mashrafi}@i2r.a-star.edu.sg

This chapter's objective is, as a survey paper, to define the domain of security data mining by organisations using published case studies from various security environments. Although each security environment may have its own unique requirements, this chapter argues that they share similar principles to operate well.

This chapter's main contribution is the focus on ways to engineer tamper-resistance for security data mining applications - mathematical algorithms in computer programs which perform security data mining. With tamper-resistance, organisations applying security data mining can better achieve accuracy for organisations, privacy for individuals in the data, and confidentiality between organisations which share the results.

This chapter is written for the general audience who has little theoretical background in data mining, but interested in practical aspects of security data mining. We assume that the reader knows about or will eventually read up on the data mining process [20] which involves ordered and interdependent steps. These steps consist of data pre-processing, integration, selection, and transformation; use of common data mining algorithms (such as classification, clustering, and association rules); results measurement and interpretation.

The rest of this chapter is organised as follows. We present security data mining's definitions, specific and general issues in Section 7.2. We discuss tamper-resistance in the form of reliable data, anomaly detection algorithms, and privacy and confidentiality preserving results in Section 7.3. We conclude with a summary and future work in Section 7.4.

7.2 Security Data Mining

This section defines terms, presents specfic as well as general problems to security data mining, and offers solutions in the form of successful applications from various security environments.

7.2.1 Definitions

The following definitions (in bold font), which might be highly evident to some readers, are specific to security data mining. An **adversary** is a malicious individual whose aim is to inflict adverse consequences to valuable assets without being discovered. Alternatively, an adversary can be an organisation, and have access to their own data and algorithms. An adversary can create more automated software and/or use more manual means to carry out an attack. Using the relevant, new, and interesting domain of virtual gaming worlds, cheating can be in the form of automated of gold farming. In constrast, cheating can also come in the form of cheap manual labour who game in teams to slaughter high-reward monsters [28].

Internal adversaries work for the organisation, such as employees responsible for data breaches [29, 46]. External adversaries do not have any access rights to the organisation, such as taxpayers who evade tax [13]. Data leak detection uses matching of documents using dictionaries of common terms and keywords, and using fingerprints of sensitive documents, and monitoring locations where sensitive documents are kept. One-class Support Vector Machines (SVM) are trained to rank new taxpayers on known-fraudulent individual and high income taxpayers data. Subsequently, the taxpayers will then be subjected to link analysis using personal data to locate pass-through entities.

Security is the condition of being protected against danger or loss. But a more precise definition of security here is the use of countermeasures to prevent the deliberate and unwarranted behaviour of adversaries [41].

Security data mining, a form of countermeasure, is the use of large-scale data analytics to dynamically detect a small number of adversaries who are constantly changing. It encompasses data- and results-related safeguards. Security data mining is relevant across multiple domains such as financial, insurance, health, taxation, social security, e-commerce, just to name a few. It is a collective term for detection of fraud, crime, terrorism, financial crime, spam, and network intrusion [37]. In addition, there are other forms of adversarial activity such as detection of online gaming [28], data breaches, phishing, and plagarism. The difference between security data mining and fraud data mining is that the former concentrates in the long-term on the adversary, not for short-term profit.

To understand security data mining better, security data mining is compared with database marketing - its opposite domain. A casino can use both domains to increase profit: Non-Obvious Relationship Awareness (NORA) [26] reduces cost, while HARRAH's database marketing [32] increases revenue. In real-time, NORA detects people who are morphing identities. NORA evaluates similarities and differences between people or organisations and shows how current entities are connected to all previous entities. In retrospect, HARRAH cultivates lasting relationships with its core customers. HARRAH discovered that slot players who are retirees are their core customers, and direct resources to develop better customer satisfaction with them.

7.2.2 Specific Issues

The following concepts (in bold font) are specific to security data mining:

- **Resilience**, for security systems, is the ability to degrade gracefully when under most real attacks. The security system needs "defence-in-depth" with multiple, sequential, and independent layers of defence [41] to cover different types of attacks, and to eliminate clearly legitimate examples [24]. In other words, any attack has to pass every layer of defence without being detected.

The security system is a combination of manual approaches; and automated approaches including blacklist matching and security data mining algorithms. The basic automated approaches include hard-coded rules such as matching personal name and address, and setting price and amount limits.

One common automated approach is known fraud matching. Known frauds are usually recorded in a periodically updated blacklist. Subsequently, the current claims/applications/transactions/accounts/sequences are matched against the blacklist. This has the benefit and clarity of hindsight because patterns often repeat themselves. However, there are two main problems in using known frauds. First, they are untimely due to long time delays which provides a window of opportunity for fraudsters. Second, recording of frauds is highly manual.

- **Adaptivity**, for security data mining algorithms, accounts for morphing fraud behaviour, as the attempt to observe fraud changes its behaviour. But what is not obvious, but equally important, is the need to also account for legal (or legitimate) behaviour within a changing environment.

  In practice, for telecommunications superimposed fraud detection [19], there is fraud rule generation from each cloned phone account's labelled data and rule selection to cover most accounts. For anomaly detection, each selected fraud rule is applied in the form of monitors (number and duration of calls) to the daily legitimate usage of each account. StackGuard [8] is a simple compiler which virtually eliminates buffer overflow attacks with only modest speed penalties. To provide an adaptive response to intrusions, StackGuard switches between the more effective MemGuard version and the more efficient Canary version.

  In theory, in spam detection, adversaries learn how to generate more false negatives from prior knowledge, observation, and experimentation [33]. Game theory is adapted to automatically re-learn a cost-sensitive supervised algorithm given the cost-sensitive adversary's optimal strategy [11]. It defines the adversary and classifier optimal strategy by making some valid assumptions.

- **Quality data** is essential for security data mining algorithms through the removal of data errors (or noise). HESPERUS [38] filters duplicates which have been re-entered due to human error or for other reasons. It also removes redundant attributes which have many missing values, and other issues. Data preprocessing for securities fraud detection [18] include known consolidation and link formation techniques to associate people with office locations, infer associations by employment histories, and normalisation techniques by space and time to create a suitable class labels.

7.2.3 General Issues

The following concepts (in bold font) are general to data mining, and are used here to describe security data mining applications.

- **Personal data versus behaviourial data**
  Personal data relates to an identified natural people, on the other hand, behavioural data relates to the actions of people under specified circumstances. The data here refers to text form, as image and video data are beyond our scope. Most applications use behavioural data but some, such as HESPERUS [38], use personal data.

  HESPERUS discovers credit card application fraud patterns. It detects sudden and sharp spikes in duplicates within a short time, relative to normal behaviour.

- **Unstructured data versus structured data**
  Unstructured data is not in a tabular or delimited format; while structured data is segmented into attributes where each has an assigned format. In this chapter's subsequent applications, most use structured data but some, such as in software plagarism [40], use unstructured data.

  Unstructured data is transformed into fingerprints - selected and hashed k-grams (using 0 mod p or winnowing) with positional information - to detect software copies. Some issues discussed in the paper include support for a variety of input formats, filter of unnecessary code, and presentation of results.

- **Real-time versus retrospective application**
  A real-time application processes events as they happen, and need to scale up to the arrival and growth of data. In contrast, a retrospective application processes events after they have taken place, and are often used to perform audits and stress tests. A real-time financial crime detection application - Securities Observation, News Analysis, and Regulation (SONAR) [22], and a retrospective managerial fraud detection application - SHERLOCK [5] are described in detail below.

  In real-time, SONAR monitors main stock markets for insider trading by using privileged information of a material nature, and misrepresentation fraud by fabricating news. SONAR mines for explicit and implicit relationships among the entities and events, using text mining, statistical regression, rule-based inference, uncertainty, and fuzzy matching.

  In retrospect, SHERLOCK analyses the general ledger - a formal listing of journal accounts in a business used for financial statement preparation and tax filing - for irregularites which are useful to auditors and investigators. SHERLOCK extracts a few dozen important attributes for outlier detection and classification. Some limitations stated in the paper include data which is hard to pre-process, having a small set of known fraud general ledgers while the major-

ity are unlabelled, and results are hard to interpret.

- **Unsupervised versus supervised application**
  An unsupervised application do not use class labels - usually assignment of records to a particular category - and is more suited for real-time use. A supervised application use class labels and is usually for retrospective use. The following click fraud detection [34] and management fraud detection [47] applications use behavioural, structured data.

  Using user click data on web advertisements, [34] analyses requests using an unsupervised pair-wise analysis with association rules. Using public company account data, use a supervised decision tree to classify time and peer attributes, and apply supervised logistic regression for each leaf time series.

- **Maximum versus no user interaction**
  Maximum user (or domain expert) interaction is required if the consequences for a security breach is severe [25]. User interaction refers to being able to easily annotate, add attributes, or change attribute weights; or to allow better understanding and use of scores (or rules). No user interaction refers to a fully automated application.

  Visual telecommunications fraud detection [9] combines user detection with computer programs. It flexibly encodes data using colour, position, size and other visual characteristics with multiple different views and levels.

7.3 Tamper-Resistance

Figure 7.1 gives a visual overview of tamper-resistance solutions in security data mining. The problems come from data adversaries, internal adversaries, and external adversaries in the form of other organisations sharing the data or results (for example, adversaries always try to look legitimate). The solutions can be summarised as **tamper-resistance**, which addresses most kinds of adversaries and makes it more difficult for an adversary to manipulate or circumvent security data mining. From experience, we recommend reliable data as inputs, anomaly detection algorithms as processes, and privacy and confidentiality preserving results as outputs to enhance tamper-resistance; and we elaborate more on them in the following subsections.

7.3.1 Reliable Data

Reliable data is not just quality data (see previous subsection 7.2.2); but also can be trusted and gives the same results, even with adversary manipulation. By reliable data, we refer to unforgeable, stable, and non-obvious data [43]. To an adversary,

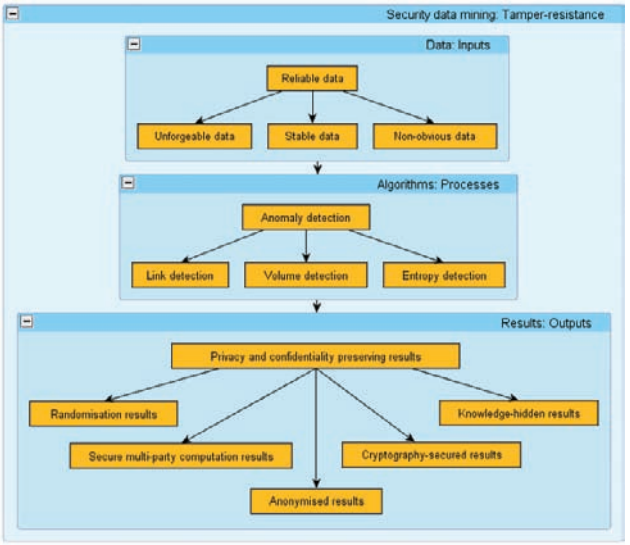

Fig. 7.1 Visual overview

reliable data cannot be replicated with the intent to deceive, has little fluctuation, and is hard to see and understand.

- **Unforgeable data** can be viewed as attributes which are generated subconsciously, such as rhythm-based typing patterns [36] which is based on timing information of username and password. As an authentication factor, rhythm-based typing patterns is cheap, easily accepted by most users, and can be used beyond keyboards. However, there exists policy and privacy issues.

- **Stable data** include communication links between adversaries where links are already available. By linking mobile phone accounts using call quantity and durations to form Communities Of Interest (COI), two distinctive characteristics of fraudsters can be determined. Fraudulent phone accounts are linked as fraudsters call each other or the same phone numbers, and fraudulent call behaviour from known frauds are reflected in some new phone accounts [7].

  Also, stable data can come from attribute extraction where attributes are not directly available or when there are too many attributes. To find new, previously unseen, malicious executables and differentiate them from benign programs, there is attribute extraction of various information such as Dynamically Linked Library (DLL) calls, consecutive printable strings, and byte sequences [42].

- **Non-obvious data** refer to attributes with characteristic distributions. For intrusion detection, these attributes describe the network traffic, such as historical averages of source and destination Internet Protocol (IP) addresses of packets, source and destination port numbers, type of protocol, number of bytes per

pacet, and time elapsed between packets. In addition, more of such attributes are from router data, such as Central Processing Unit (CPU), memory usage, and traffic volume [35].

In online identity theft, each phishing attack has several stages starting from delivery of attack to ending of receiving of money [15]. The point here is to collect and mine non-obvious data from stages where adversaries least expect.

7.3.2 Anomaly Detection Algorithms

To detect anomalies early (also known as abnormalities, deviations, or outliers), anomaly detection algorithms are a type of security data mining algorithm which originate from network intrusion detection research [14]. They profile normal behaviour (also known as norm or baseline) by outputing suspicion scores (or rules). Anomaly detection algorithms can be used on data for various security environments, in different stages, at different levels of granularity (such as at the global, account, or individual levels), or for groups of similar things (such as dates, geography, or social groups).

Anomaly detection algorithms require class imbalanced data - plenty of normal compared to anomalous behaviour which is common in security data - to be useful [16]. They are only effective when normal behaviour has high regularity [30]. The common anomaly detection algorithms monitor for changes in links, volume, and entropy:

- **Link detection** is to find good or bad connections between things. For rhythm-based typing where links have to be discovered, [36] uses classifiers to measure the link (or similarity) between an input keystroke timing and a model of normal behaviour of the keystroke timing. Each attribute for a model of normal behaviour is an updated average from a predefined number of keystrokes. For telecommunications fraud detection where links are available, [7] examines temporal evolution of each large dynamic graph for subgraphs called COIs. For professional software plagarism detection, [31] mines Program Dependence Graphs (PDG) which links code statements based on data and control dependencies which reflects developers' thinking when code is written.

  However, for securities fraud detection, although fraud is present when there are links between groups of representatives that pass together through multiple places of employment, these links will also find harmless sets of friends that worked together in the same industry and a multitude of non-fraud links [21].

- **Volume detection** is to monitor the significant increase or decrease in amount of something. For credit card transactional fraud detection, Peer Group Analysis [6] monitors inter-account behaviour by comparison of the cumulative mean weekly amount between a target account and other similar accounts (peer

group). Break Point Analysis [6] monitors intra-account behaviour by detecting rapid weekly spending. A neural network trained on a seven day moving window of Automated Teller Machines' (ATM) daily cash output to detect anomalies. For bio-terrorism detection, Bayesian networks [48] observe sudden increases of certain illness symptoms from real emergency department data. Time series analysis [23] tracks daily sales of throat, cough, and nasal medication; and some grocery items such as facial tissues, orange juice, and soup.

However, sudden increases in volume can come from common non-malicious fluctuations.

- **Entropy detection** is to measure the sudden increase of decrease of disorder (or randomness) in a system. Entropy is a function of $k \log p$, where k is a constant and p is the probability of a given configuration. For network intrusion detection, the traffic and router attributes are characterised by distinct curves which uniquely profile the traffic: high entropy is represented by a uniform curve, while low entropy is shown as a spike.

  Even if adversaries try to look legitimate and keep usage volume low, entropy detection can still find network intrusions which differ in some way from the network's established usage patterns [30, 35].

7.3.3 Privacy and Confidentiality Preserving Results

Data sharing and data mining can be good (increases accuracy), but data mining can be bad (decreases privacy and confidentiality) [10]. For example, suppose a drug manufacturing company wishes to collect responses (i.e. record) from each of the clients containing their dining habits and adverse effects of a drug. The relationship between dining habits and the drug could give the drug manufacturing company some insight knowledge about its side effects. The clients may not be interested to provide information because of their privacy.

Another example [45] is a car manufacturing company who incorporates several components such as tires, electrical equipments, etc. made by independent producers. Each of these producers has their proprietary databases which it may not be interested to share. However, in practical scenarios sharing those databases is important and we could take the Ford Motors and Firestone Tires provide a real example of this type. Ford Explorers with Firestone tires from a specific factory had tire-thread separation problem which resulted in 800 injuries. As those tires did not cause problems to other vehicles or the other tires in Ford Explorer did not pose such problems, thus neither Ford nor Firestone wants to take responsibility. The delays in identifying the real problem resulted in public concern and eventually led to replacement of 14.1 million tires. In reality, many of those tires were probably fine as Ford Explorer accounted for only 6.5 million of the replacement tires. If both companies had discovered the association between the different attributes of their

proprietary databases, then this safety concern can be avoided before it becomes public.

Privacy and confidentiality are important issues in security data mining because organisations use personal, behavioural, and sensitive data. Explicit consent has been given by the people to use their personal data and behaviourial data for a specific purpose, and all personal data is protected from unauthorised disclosure or intelligible interception. Privacy laws, non-disclosure agreements, and ethical codes of conduct have to be adhered to. Sometimes, the exchange of raw or summarised results with other organisations may expose personal or sensitive data. Therefore, the following are ways to increase privacy and confidentiality, mainly from association rules literature:

- **Randomisation** is simple probabilistic distortion of user data, employing random numbers generated from a pre-defined distributed function. A centralised environment to maintain privacy and accuracy of resultant rules has been proposed [39]. However, the distortion process employs system resources for a long period when the dataset has large number of transactions. Furthermore, if this algorithm is used in the context of a distributed environment, this needs uniform distortion among various sites in order to generate unambiguous rules. This uniform distortion may disclose confidential inputs of individual site and may also breach the privacy of data (such as exact support of itemsets), and hence it is not suitable for large distributed data mining.

  To discover patterns from distributed datasets, a randomisation technique [17] could be deployed in an environment where a number of clients are connected to a server. Each client sends a set of items to the server where association rules are generated. During the sending process, the client modifies the itemsets according to its own randomisation policies. As a result, the server is unable to find the exact information about the client.

  However, this assumption is not suitable for distributed association rule mining because it generates global frequent itemsets by aggregating support counts of all clients.

- **Secure Multi-party Computation (SMC)**-based [3,27] perform a secure computation at individual site. To discover a global model, those algorithms secretly exchange the statistical measures among themselves. This is more suitable for few external parties.

  A privacy preserving association rule mining is defined for horizontally partitioned data (each site shares a common schema but has different records) [27]. Two different protocols were proposed: secure union of locally large itemsets, and a testing support threshold without revealing support counts. The former protocol uses cryptography to encrypt local support count, and therefore, it is not possible to find which itemset belongs to which site. However, it reveals the local itemsets to all participating sites in where these itemset are also locally frequent. Since the first protocol gives the full set of locally frequent itemsets, then in order to find which of these itemsets are globally frequent, the latter

protocol is used. It adds a random number to each support count and finds the excess supports. Finally, these excess supports are sent to the second site where it learns nothing about the first site's actual dataset size or support. The second site adds its excess support and sends the value until it reaches the last site.

This protocol can raise a collusion problem. For example, site i and $i+2$ in the chain can collude to find the exact excess support of site $i+1$. To generate patterns from vertically partitioned distributed dataset, a technique is used to maintain the privacy of resultant patterns in vertically partitioned distributed data sources (across two data sources only) [45]. Each of the parties holds some attributes of each transaction. However, if the number of disjoint attributes among the site is high, this technique incurs huge communication costs. Furthermore, this technique is designed for an environment where there are two collaborating sites, each of them holding some attributes of each transaction. Hence, it may not be applied in an environment where collaborating sites do not possess such characteristics.

- **Anonymity** minimises potential privacy breaches. The above two techniques - randomisation and SMC - focused on how to find frequency of itemsets from large dataset such a way that none of the participants is able to see the exact local frequency each of the individual itemset. Though the patterns discovered using these methods does not reveal exact frequency of an itemset however, the resultant patterns may reveal some information about the original dataset which are not intentionally released. In fact, such inferences represent *per se* a threat to privacy. To overcome such potential threat, k-anonymous patter discovery method is proposed [4]. Unlike the randomisation, the proposed method generates patterns using data mining algorithm from the real dataset. Then these patterns are analysed against several anonymity properties.

  The anonymity properties check whether collection of patterns guarantee the anonymity or not. Based on the outcome of the anonymity, the patterns collection is sanitised in such a way that the anonymity of a given pattern collection is preserved. As the patterns are generated using the real dataset, the main problem of this approach is how to discover patterns from distributed datasets. In fact, if each of the participating sites applies this method at local sites, then the resultant global patterns will have discrepancies which could diminish the goal of distributed pattern mining.

- **Cryptography**-based techniques [27, 49] use the public key cryptography system to generate a global model. This is more suitable for many external parties. Despite cryptography system has computational and communication overhead, recent research argues it is possible to generate privacy preserving patterns and with achieve good performance. For example, a cryptography-based system that performs sufficiently efficient to be useful in the practical data mining scenarios [49]. Their proposed method discovers patterns in a setting where number of participant is large. Each of the participants sends their own private input to a

data miner who will generate patterns from these inputs using the homomorphic property of ElGamal cryptography system.

The main problem of cryptography-based approach is the underlying assumptions. For example, all of the cryptography-based methods assume participating parties are semi-honest, that is, each of them executes the protocol exactly the same manner as described in the protocol specification. Unless each of the participants is semi-honest, those methods may not able to preserve the privacy of the each of the participant's private input.

- **Knowledge hiding** is a way to preserve privacy of sensitive knowledge by hiding frequent itemsets from large datasets. Heuristics were applied to reduce the number of occurrences to such a degree that its support is below the user-specified support threshold [2]. This work was extended to confidentiality issues of association rule mining [12]. Both works assume datasets are local and that hiding some itemsets will not affect the overall performance or mining accuracy.

  However, in distributed association rule mining, each site has its own dataset and a similar kind of assumption may cause ambiguities in the resultant global rule model.

7.4 Conclusion

This chapter is titled *Security Data Mining: A Survey Introducing Tamper-Resistance*, that is, motivations, definitions, and problems are discussed and tamper-resistance as an important solution is recommended. The growth of security data with adversaries has to be accompanied by both theory-driven and domain-driven data mining. Inevitably, security data mining with tamper-resistance has to incorporate domain-driven enhancements in the form of reliable data, anomaly detection algorithms, and privacy and confidentiality preserving results. Future work will be to apply tamper-resistance solutions to the detection of data breaches, phishing, and plagarism; for specific results to support the conclusion of this chapter.

References

1. Adams, N.: 'Fraud Detection in Consumer Credit'. Proc. of UK KDD Workshop (2006)
2. Atallah, M., Bertino, E., Elmagarmid, A., Ibrahim, M., Verykios, V.,: 'Disclosure Limitation of Sensitive Rules'. Proc. of KDEX99, pp. 45–52 (1999)
3. Ashrafi, M., Taniar, D., Smith, K.: 'Reducing Communication Cost in a Privacy Preserving Distributed Association Rule Mining'. Proc. of DASFAA04, LNCS 2973, pp. 381–392 (2004)
4. Atzori, M., Bonchi, F., Giannotti, F., Pedreschi, D.: '*k*-Anonymous Patterns'. Proc. of PKDD05, pp. 10–21 (2005)
5. Bay, S., Kumaraswamy, K., Anderle, M., Kumar, R., Steier, D: 'Large Scale Detection of Irregularities in Accounting Data'. Proc. of ICDM06, pp. 75–86 (2006)

6. Bolton, R., Hand, D.: 'Unsupervised Profiling Methods for Fraud Detection'. Proc. of CSCC01 (2001)
7. Cortes, C., Pregibon, D., Volinsky, C.: 'Communities of Interest'. Proc. of IDA01. pp. 105–114 (2001)
8. Cowan, C., Pu, C., Maier, D., Walpole, J., Bakke, P., Beattie, S., Grier, A., Wagle, P., Zhang, Q., Hilton, H: 'StackGuard: Automatic Adaptive Detection and Prevention of Buffer-Overflow Attacks'. Proc. of 7th USENIX Security Symposium (1998)
9. Cox, K., Eick, S., Wills, G.: 'Visual Data Mining: Recognising Telephone Calling Fraud'. *Data Mining and Knowledge Discovery* **1**. pp. 225–231 (1997)
10. Clifton, C., Marks, D.: 'Security and Privacy Implications of Data Mining'. Proc. of SIGMOD Workshop on Data Mining and Knowledge Discovery. pp. 15–19 (1996)
11. Dalvi, N., Domingos, P., Mausam, Sanghai, S., Verma, D.: 'Adversarial Classification'. Proc. of SIGKDD04 (2004)
12. Dasseni, E., Verykios, V., Elmagarmid, A., Bertino, E.: 'Hiding Association Rules by Using Confidence and Support'. LNCS 2137, pp. 369-379 (2001)
13. DeBarr, D., Eyler-Walker, Z.: 'Closing the Gap: Automated Screening of Tax Returns to Identify Egregious Tax Shelters'. *SIGKDD Explorations*. **8**(1), pp. 11–16 (2006)
14. Denning, D.: 'An Intrusion-Detection Model'. *IEEE Transactions on Software Engineering*. **13**(2), pp. 222–232 (1987)
15. Emigh, A., 'Online Identity Theft: Phishing Technology, Chokepoints and Countermeasures'. ITTC Report on Online Identity Theft Technology and Countermeasures (2005)
16. Eskin, E., Arnold, A., Prerau, M., Portnoy, L., Stolfo, S.: 'A Geometric Framework for Unsupervised Anomaly Detection: Detecting Intrusions in Unlabeled Data'. Applications of Data Mining in Computer Security, Kluwer (2002)
17. Evfimievski, A., Srikant, R., Agrawal, R., Gehrke, J.: 'Privacy Preserving Mining of Association Rules', *Information Systems*, **29**(4): pp. 343–364 (2004)
18. Fast, A., Friedland, L., Maier, M., Taylor, B., Jensen, D., Goldberg, H., Komoroske, J.: 'Relational Data Pre-Processing Techniques for Improved Securities Fraud Detection'. Proc. of SIGKDD07 (2007)
19. Fawcett, T., Provost, F.: 'Adaptive Fraud Detection'. *Data Mining and Knowledge Discovery*. **1**(3), pp. 291–316 (1997)
20. Fayyad, U., Piatetsky-Shapiro, G., Smyth, P., Uthurusamy, R.: Advances in Knowledge Discovery and Data Mining. AAAI (1996)
21. Friedland, L., Jensen, D.: 'Finding Tribes: Identifying Close-Knit Individuals from Employment Patterns'. Proc. of SIGKDD07 (2007)
22. Goldberg, H., Kirkland, J., Lee, D., Shyr, P., Thakker, D: 'The NASD Securities Observation, News Analysis and Regulation System (SONAR)'. Proc. of IAAI03 (2007)
23. Goldenberg, A., Shmueli, G., Caruana, R.: 'Using Grocery Sales Data for the Detection of Bio-Terrorist Attacks'. *Statistical Medicine* (2002)
24. Hand, D.: 'Protection or Privacy? Data Mining and Personal Data'. Proc. of PAKDD06, LNAI 3918. pp. 1–10 (2006)
25. Jensen, D.: 'Prospective Assessment of AI Technologies for Fraud Detection: A Case Study'. AI Approaches to Fraud Detection and Risk Management. AAAI Press, pp. 34–38 (1997)
26. Jonas, J.: 'Non-Obvious Relationship Awareness (NORA)'. Proc. of Identity Mashup (2006)
27. Kantarcioglu, M., Clifton, C.: 'Privacy-Preserving Distributed Mining of Association Rules on Horizontally Partitioned Data'. *IEEE Transactions on Knowledge and Data Engineering*. **16**(9), pp. 1026–1037 (2004)
28. Kushner, D.: 'Playing Dirty: Automating Computer Game Play Takes Cheating to a New and Profitable Level'. *IEEE Spectrum.* **44**(12) (INT), December 2007, pp. 31–35 (2007)
29. Layland, R.: 'Data Leak Prevention: Coming Soon To A Business Near You'. *Business Communications Review*. pp. 44–49, May (2007)
30. Lee, W., Xiang, D.: 'Information-theoretic Measures for Anomaly Detection'. Proc. of 2001 IEEE Symposium on Security and Privacy (2001)
31. Liu, C., Chen, C., Han, J., Yu, P.: 'GPLAG: Detection of Software Plagiarism by Program Dependence Graph Analysis'. Proc. of SIGKDD06 (2006)

32. Loveman, G.: 'Diamonds in the Data Mine'. *Harvard Business Review*. pp. 109–113, May (2003)
33. Lowd, D., Meek, C.: 'Adversarial Learning'. Proc. of SIGKDD05 (2005)
34. Metwally, A., Agrawal, D., Abbadi, A.: 'Using Association Rules for Fraud Detection in Web Advertising Networks'. Proc. of VLDB05 (2005)
35. Nucci, A., Bannerman, S.: 'Controlled Chaos'. *IEEE Spectrum*. **44**(12) (INT), December 2007, pp. 37–42 (2007)
36. Peacock, A., Ke X., Wilkerson, M.: 'Typing Patterns: A Key to User Identification'. *IEEE Security and Privacy* **2**(5), pp. 40–47 (2004)
37. Phua, C., Lee, V., Smith-Miles, K., Gayler, R.: 'A Comprehensive Survey of Data Mining-based Fraud Detection Research'. Clayton School of Information Technology, Monash University (2005)
38. Phua, C.: 'Data Mining in Resilient Identity Crime Detection'. PhD Dissertation, Monash University (2007)
39. Rizvi, S., Haritsa, J.: 'Maintaining Data Privacy in Association Rule Mining'. Proc. of VLDB02 (2002)
40. Schleimer, S., Wilkerson, D., Aiken, A.: 'Winnowing: Local Algorithms for Document Fingerprinting'. Proc. of SIGMOD03. pp. 76–85 (2003)
41. Schneier, B.: Beyond Fear: Thinking Sensibly about Security in an Uncertain World. Copernicus (2003)
42. Schultz, M., Eskin, E., Zadok, E., Stolfo, S.: 'Data Mining Methods for Detection of New Malicious Executables'. Proc. of IEEE Symposium on Security and Privacy. pp. 178–184 (2001)
43. Skillicorn, D.: Knowledge Discovery for Counterterrorism and Law Enforcement. CRC Press, in press (2008)
44. Sweeney, L.: 'Privacy-Preserving Surveillance using Databases from Daily Life'. *IEEE Intelligent Systems*. **20**(5): pp. 83—84 (2005)
45. Vaidya, J., Clifton C.: 'Privacy Preserving Association Rule Mining in Vertically Partitioned Data'. Proc. of SIGKDD02.
46. Viega, J.: 'Closing the Data Leakage Tap'. *Sage*. **1**(2): Article 7, April (2007)
47. Virdhagriswaran, S., Dakin, G.: 'Camouflaged Fraud Detection in Domains with Complex Relationships'. Proc. of SIGKDD06 (2006)
48. Wong, W., Moore, A., Cooper, G., Wagner, M.: 'Bayesian Network Anomaly Pattern Detection for Detecting Disease Outbreaks'. Proc. of ICML03 (2003)
49. Yang, Z., Zhong, S., Wright, R.: 'Privacy-Preserving Classification of Customer Data without Loss of Accuracy'. Proc. of SDM05 (2005)