



# AI and Security Informatics

Hsinchun Chen, University of Arizona

Since the tragic events of September 11, security research has become critically important for the entire world. Academics in fields such as computational science, information systems, social sciences, engineering, and medicine have been called on to help enhance our ability to fight violence, terrorism, and other crimes. The US 2002 National Strategy for Homeland Security report identified science and technology as the keys to winning this international security war.<sup>1</sup> It is widely believed that information technology will play an indispensable role in making the world safer<sup>2</sup> by supporting intelligence and knowledge discovery through collecting, processing, analyzing, and utilizing terrorism- and crime-related data.<sup>3</sup>

Based on the available crime and intelligence knowledge, federal, state, and local authorities can make timely and accurate decisions to select effective strategies and tactics as well as allocate the appropriate amount of resources to detect, prevent, and respond to future attacks. Facing the critical mission of international security and various data and technical challenges, there is a pressing need to develop the science of *security informatics*. The main objective is the development of advanced information technologies, systems, algorithms, and databases for security-related applications using an integrated technological, organizational, and policy-based approach. Intelligent systems have much to contribute for this emerging field.

## The Dark Web Forum Portal

In recent years, there have been numerous studies from various perspectives analyzing the Internet presence of extremist and terrorist groups. Yet the websites and forums of extremist and terrorist groups have long remained an underutilized

resource for terrorism researchers due to their ephemeral nature as well as access and analysis problems. To address this gap, the University of Arizona Artificial Intelligence Lab's Dark Web archive (<http://ai.eller.arizona.edu/research/terror/>) provides a research infrastructure for use by social scientists, computer and information scientists, policy and security analysts, and others studying a range of social and organizational phenomena and computational problems.

The Dark Web archive currently consists of 13 million postings from 29 international jihadist Web forums. These forums collectively host 340,000 members, whose discussions cover a range of socio-political, cultural, ideological, and religious topics. The forums collected are in Arabic, English, French, German, and Russian and have been carefully selected with significant input from terrorism researchers, security and military educators, and other experts.

The Dark Web Forum Portal system currently consists of four types of functions:

- single and multiple forum browsing and searching,
- forum statistics analysis,
- multilingual translation, and
- social network visualization.

The search function lets users search message titles or bodies using multiple keywords. In addition to browsing and searching information in a particular forum, the portal also supports searching across all the forums. Forum statistics are also carefully summarized. All search terms and message postings are translated automatically using Google Translate (<http://translate.google.com>). The system also supports forum participant network visualization, using selected social network

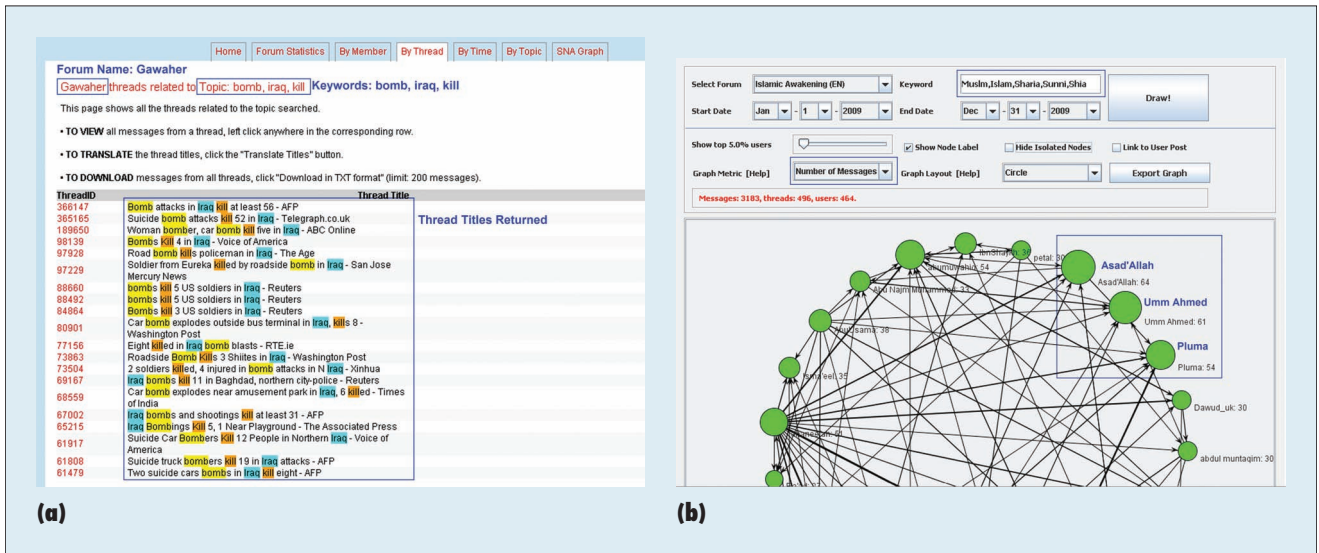


Figure 1. Dark Web Forum Portal interface. The portal (a) allows searches across all the forums and (b) displays forum participant network search and visualization.

analysis (SNA) metrics and visualization methods.

Figure 1 shows a sample search screen and SNA visualization. Funded partially by the National Science Foundation’s Computation Research Infrastructure (CRI) program, the Dark Web Forum Portal is a scalable infrastructure that integrates heterogeneous forum data and serves as a strong complement to existing security databases, news reports, and other sources available to the security informatics research community.

### In This Issue

This issue includes three articles on AI and security informatics from distinguished experts in computer science and information systems. Each article presents a unique, innovative research framework, computational methods, and selected results and examples.

In “Building Models When Adversary Data Is Rare,” David B. Skillicorn argues that security and intelligence data often contain a large number of records associated with good actions, but few, perhaps none, associated with bad actions. He proposes several models to address this issue.

Bhavani Thuraisingham, Latifur Khan, and Murat Kantarcioglu’s article “Semantic Web, Data Mining and Security” examines the integration of security technologies with the Semantic Web and data mining information technologies. They discuss how Semantic Web and security technologies might be integrated, including a discussion of inference control and policy specifications. They also describe how to integrate data mining and security technologies, including a discussion of privacy-preserving data mining techniques.

Lastly, in the “Privacy-Preserved Social Network Integration and Analysis for Security Informatics,” Christopher C. Yang and Bhavani Thuraisingham discuss the strength and limitations of leading approaches for privacy-preserving data mining based on anonymization models. They introduce the subgraph generalization approach for social network integration and demonstrate its feasibility for integrating social networks and preserving privacy.

### Acknowledgments

This material is based on work supported in part by the following grants: Department

of Defense HDTRA-09-0058, National Science Foundation CNS-070933, NSF CBET-0730908, and NSF IIS-0428241.


### References

1. “National Strategy for Homeland Security,” US Office of Homeland Security, 2002.
2. Nat’l Research Council, “Making the Nation Safer: The Role of Science and Technology in Countering Terrorism,” Nat’l Academy Press, 2002.
3. H. Chen, *Intelligence and Security Informatics for International Security: Information Sharing and Data Mining*, Springer, 2006.

### Building Models When Adversary Data Is Rare

David B. Skillicorn, *Queen’s University*

Security and intelligence data often contains numerous of records associated with good actions, but few, perhaps none, associated with bad actions. For example, aircraft passenger screening systems contain billions of records about normal travelers for each record about a terrorist or hijacker. In crime, fraud, tax evasion, and money laundering, the



**Figure 2.** A simplified model of good (blue) and bad (red) records and a possible boundary between them. The dashed circle indicates a two-class predictor that learns the boundary between the classes.

ratio of good to bad records might not be quite as great, but examples of bad records are still rare. Many of these settings are not only unbalanced in the number of examples of different kinds of records, they are also adversarial because those whose actions generate the bad records are trying to conceal themselves and also perhaps disrupt the modeling that is being done. In these settings, we still need to be able to build robust inductive models (predictors, clusterings, and rankings) from the data, but there are new issues to be considered.

### Separating Good and Bad Records

Figure 2 is a simplified illustration of data from such a setting, with the good records shown as stars, and bad records as circles. There are clearly differences between good and bad records, but how can we robustly model these differences? One approach is to build a two-class predictor that learns the boundary between the classes, perhaps using a support vector machine with a radial basis function kernel, indicated by the dashed circle in Figure 2.

This solution has several limitations. First, there are not enough bad points to allow generalization to define a bad region; the potential

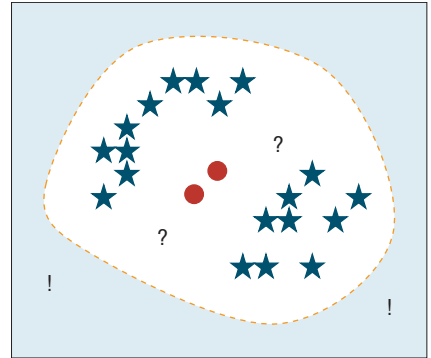


**Figure 3.** An example of a one-class classification. Anything outside the dotted lines is bad.

bad region might not be big enough in some dimensions (undergeneralization) and too big in others. Second, the good points define a region that is much larger than the region they actually occupy because it includes everything outside the dashed circle—the potential good region is overgeneralized. The predictions of such a model will probably include large numbers of both false positives and false negatives. In practice, this will make the model useless, a point already observed in previous research.<sup>1</sup>

A second approach is to pose the problem as one-class prediction<sup>2</sup>—generalize the good records to a good region, and treat anything outside this region as bad. To do this, no bad records are actually necessary; bad is defined as “not good,” as Figure 3 illustrates.

There are several ways to model the good region, but they all wrap the good records as a way of generalizing what these are like. For example, the good region could be the convex hull of the good records, a solution that implicitly assumes that records with attributes that are linear combinations of good record attributes are probably also good. If the distribution of good records is less compact, the good region could be the union of small regions centered on each good record—for example, spheres or cubes.



**Figure 4.** New records are either novel (the exclamation marks outside the dashed region) or interesting (the question marks inside the dashed region but not close to known records).

This approach produces very different results from the first; any record that does not closely resemble a good record will be considered a bad record. There will be very few false negatives, but the number of false positives will almost certainly be larger.

### Records that Don't Seem Either Good or Bad

The problem with both of these approaches is that they classify new records as good or bad, even if they are not at all similar to those that were used to build the model.<sup>3</sup> This kind of unjustified generalization is problematic in adversarial settings (indeed, in many settings). It is better to admit that there is never enough information to model every possible new record unambiguously and build that into the process.

In fact, the problematic records can be thought of as of two different kinds. The first kind is those that do not resemble either the known good or bad records. In Figure 4, these *novel* records lie outside the dashed line and are indicated by exclamation marks. Their presence is a signal that the existing model is inadequate in its coverage.

The second kind is those that lie between existing good and bad records,

Table 1. Models of expectation.

Model	Purpose	Exception meaning
Input model	Detects records unlike those that used to build the main model	New record is novel, unlike data used to build the model
Mapping model	Detects records for which the main model mapping is poor	New record is interesting, existing model is inadequate
Change model	Detects records and mappings that are changing systematically over time	Model is becoming outdated

but do not seem similar to either. In Figure 4, these *interesting* records are indicated by question marks. Their presence is a signal that the existing model is inadequate in its richness; perhaps boundaries are misplaced or the number of apparent clusters is too small or too large. In adversarial settings, interesting records might also be considered suspicious; the opposing pressures of bad activity and a desire for concealment often lead to records that are nearly, but not quite, normal.

It is also important, especially in adversarial settings, to be aware of the potential for manipulation and the existence of a kind of arms race between modelers and those who wish to remain hidden or invisible. This requires awareness that any model is a temporary representation of the system being modeled and that models might have to be rebuilt, perhaps often.

### Models of Expectation

The core of intelligence, in the ordinary sense, is the existence of a model of what should happen. We compare this expectation model with what does happen, the reality, and the difference provides the mechanism for “going meta” that prevents, for example, humans from getting stuck in the way software sometimes does.

All inductive data modeling, but especially in adversarial settings, should include such a model of what should happen, supplementing the main model. In fact, three models are appropriate (see Table 1.)

The first supplementary model describes what input should be like. It models the region of previously-seen records, for example, the region inside the dashed line in Figure 4. New records that do not fit this input model are treated as novel. Even in

situations where the range of possible new records is extremely large (for example, in intrusion detection), there is often a strong prior that provides extra information about whether a never-before-seen record is likely to be good or bad, so the input model could produce a continuous output—that is, the probability of a new record being novel, rather than a categorical output.

The second supplementary model describes what input mapping should be like. This mapping model could be a simpler version of the main model or could use some extra information from the main model. For example, in ensemble models, the margin of the winning vote is a strong indication of the model’s confidence in its prediction, a surrogate for how much the new record resembles the model’s current structure. New records that fit the input model but do not fit the mapping model are labeled interesting.

The third supplementary model describes how stable the modeling process is, perhaps by watching the distribution of incoming new records, the distribution of mappings, and the error rates. This change model watches for disagreement between the modeling process and the real world.

What happens when a new record does not fit well with one of these models? There are many choices, but the point is that there are choices because the new record’s unusual properties have been noticed and would not have been in a more conventional

process. In other words, failing to fit one of the expectation models provides a hook for metaprocessing, just as it does for humans. Such a record could be passed to a more sophisticated process that can afford to be more comprehensive because it is invoked less often. This might result in more sophisticated modeling, capture of, and access to extra (more expensive) attributes<sup>4</sup> or even human intervention. When the change model indicates that the outside environment is changing, it is time to rebuild the model.

### An Example: Spam

The potential use of these models can be seen in the familiar domain of email spam filtering. Novel emails are those containing words never seen before. Current spam-detection systems classify these as non-spam, although spammers exploit this by creating look-alike versions of words likely to be flagged. One possibility created by using an input model is that novel emails could be diverted and analyzed for the presence of unusual spelling or punctuation symbols within words. Interesting emails are those with a low confidence in the classification as spam or non-spam. Current systems pass emails that are weakly considered non-spam to the user, but not those that are weakly considered spam. Passing through a few of these latter emails and flagging both as uncertain would allow improved user feedback, faster model improvement, and perhaps better effective error rates.

Finally, users typically experience periodic increases in undetected spam whenever spammers start to use some new exploit. A change model would notice this automatically and thus notify a system administrator or invoke a temporarily increased threshold or a number of other possible responses.

## References

1. J. Jonas and J. Harper, "Effective Counterterrorism and the Limited Role of Predictive Data Mining," *Policy Analysis*, vol. 584, Cato Inst., 2006, pp. 1–12.
2. D.M.J. Tax, "One Class Classification," doctoral dissertation, Technical Univ. Delft, 2000.
3. M.A.J. Bourassa and D.B. Skillicorn, "Hardening Adversarial Prediction with Anomaly Tracking," *Proc. Int'l Conf. IEEE Intelligence and Security Informatics*, IEEE Press, 2009, pp. 43–48.
4. K.A. Taipale, "Data Mining and Domestic Security: Connecting the Dots to Make Sense of Data," *Columbia Science and Technology Law Rev.*, vol. 5, no. 2, Dec. 2003.

**David B. Skillicorn** is a professor at Queen's University, Canada. Contact him at [skill@cs.queensu.ca](mailto:skill@cs.queensu.ca).

## Semantic Web, Data Mining, and Security

**Bhavani Thuraisingham, Latifur Khan, and Murat Kantarcioglu**, *University of Texas at Dallas*

Data has become such a critical resource in many organizations that efficiently accessing the data, sharing

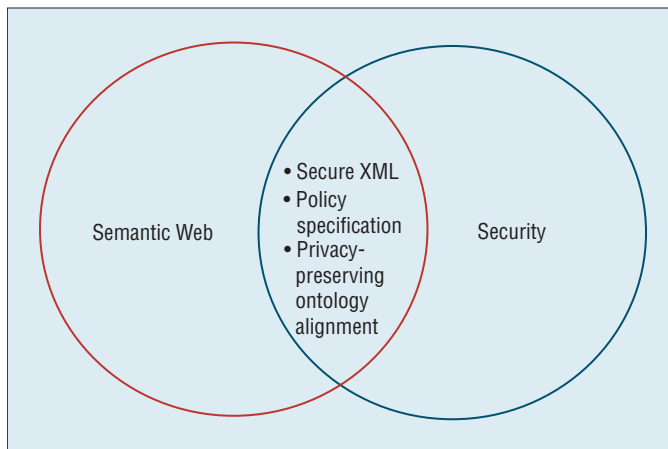


Figure 5. Intersection of Semantic Web and security.

it, extracting information from it, and making use of that information has become an urgent need. Hence, it is necessary to protect data and information from unauthorized access as well as from malicious corruption. In recent years, many efforts have researched *security informatics*, which looks at how to integrate security technologies and information technologies.<sup>1</sup>

The advent of the Web in the mid-1990s resulted in an even greater demand for effectively managing data, information, and knowledge. There is now so much data on the Web that managing it with conventional tools is becoming almost impossible. New tools and techniques are needed to effectively manage this data. Therefore, Tim Berners Lee conceived the Semantic Web to provide interoperability as well as to ensure machine-understandable webpages.<sup>2</sup> Furthermore, as information becomes unmanageable, we need data mining techniques to analyze the information and extract the nuggets often previously unknown.

Our discussion here focuses on this pressing need and how researchers might integrate security technologies with Semantic Web and data mining information technologies. Along this line, we also discuss inference control and policy specifications as well as privacy-preserving data mining techniques.

## Semantic Web and Security

Figure 5 illustrates how the Semantic Web and security technologies can benefit from one another. Languages such as XML and the Resource Description Framework (RDF) are being used extensively to represent documents, but appropriate policies have to be

enforced on these documents. Research on securing XML and XML schemas is looking at how to control access to various portions of the document for reading, browsing, and modifications. To secure RDF, however, we must secure XML, which RDF is based on, and we need security for semantic interpretations. For example, we might want to classify the statements "Tim Berners Lee is the inventor of the Semantic Web" or "H.C. Chen coined the term security informatics."

Another Semantic Web security application is in the area of *ontology alignment*, which helps identify matching concepts across ontologies. We are investigating an approach based on the path difference among concepts in the ontology combined with the approximate privacy-preserving matching technique.<sup>3</sup> Path difference consists of comparing the path leading from each concept to the root of the ontology. The more similar these paths are, the more semantically similar the concepts. Because we can compactly code these paths as strings, by adopting, for example, some numeric identifiers for concepts, we can replace the corresponding values in the original records. We can then match these modified records by using the approximate privacy-preserving matching technique.<sup>3</sup> This matching technique privately

computes the distances between pairs of records and returns only the records with a distance that is below a certain threshold.

Semantic Web technologies can also be applied for security problems. The significant advantage of Semantic Web technologies is their representational and reasoning power. Due to the representational power, these technologies can be used not only for representing the data, but also the security policies. For example, XML-based languages such as Extensible Access Control Markup Language (XACML) have been used to specify various types of security policies. Researchers have proposed extensions for XACML for finer-grained access control. Furthermore, languages such as RDF are also being explored for representing the policies. Reasoning engines such as Jena and Pellet are also being explored for representing and reasoning about Semantic Web-based policy specifications and determining whether there are security violations through inference.<sup>4</sup> For example, Jena manages RDF graphs, and Pellet reasons with RDF graphs. We have applied these technologies to develop an inference controller.<sup>5</sup>

### Data Mining and Security

Figure 6 shows how data mining and security technologies can benefit from each other as well. For example, researchers are applying data

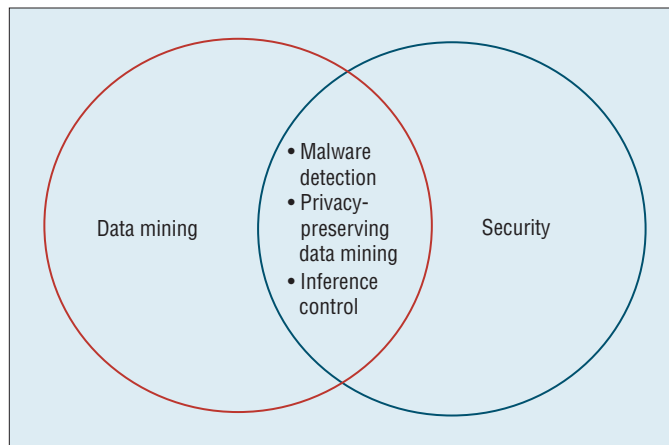


Figure 6. Intersection of data mining and security.

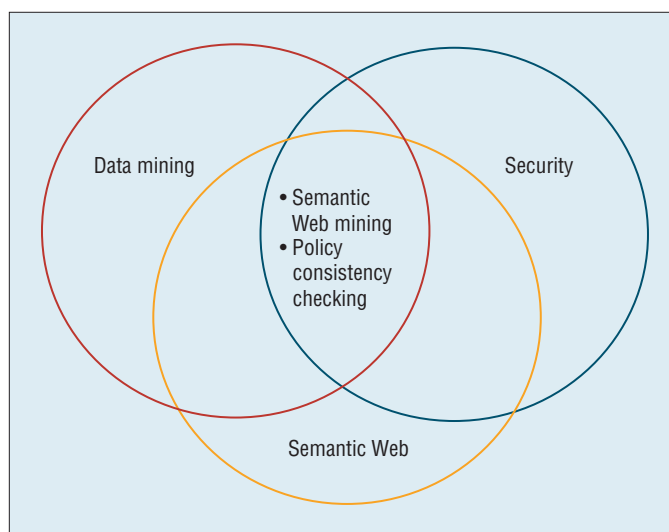


Figure 7. Intersection of Semantic Web, security, and data mining. To protect against security violations via inferences requires applying privacy-preserving data mining for XML and RDF data.

mining to problems such as intrusion detection and auditing. Anomaly detection techniques could be used to detect unusual patterns and behaviors. Link analysis could help trace viruses to the perpetrators. Classification might be used to group various cyberattacks and then use the profiles to detect an attack when it occurs. Prediction might help determine potential future attacks using information learned about terrorists through email and phone conversations. Data mining can even be used to analyze weblogs and audit trails.

Data mining can also benefit from security. First, the data mining tasks must be assured. For example, we must ensure that the data mining tasks are not maliciously corrupted. Of more concern are the inference, aggregation, and privacy problems with respect to data mining. For example, data mining gives us associations between entities that are not visible to humans but that might be highly sensitive or private. Data mining might also result in nuggets that the users can utilize to infer highly sensitive information.

With respect to privacy and confidentiality, the challenge is determining how to protect the privacy or confidentiality of the individual data while giving out results. Several privacy-preserving data mining<sup>6</sup> approaches have been proposed to address this. In one approach, the individual data values are

perturbed or random values are introduced. The goal is to mine the perturbed data but still get the original results. In another approach called *multiparty computation*,<sup>7</sup> each party knows its own inputs and the results of mining, but they do not know anything about their partners' data.

### Intersection of Semantic Web, Data Mining, and Security

Figure 7 illustrates the intersection of data mining, security, and the Semantic Web. A document might be represented in XML or RDF. To ensure

that Semantic Web mining does not result in security violations via inferences, we need to apply privacy-preserving data mining for XML and RDF data. Semantic Web technologies might also be used to specify policies. Policy consistency can be determined using data mining techniques. That is, we can apply association rule mining techniques to eliminate redundant and inconsistent policies specified in XML or RDF.

Much work is necessary to prepare the way for the integration of the Semantic Web, data mining, and security. Although some of the technologies such as privacy-preserving data mining are fairly mature, research in areas such as privacy-preserving ontology alignment is just beginning. Privacy-preserving third-party data publication and privacy-preserving information integrations are also fruitful areas of research in security informatics.

## Acknowledgments

We thank Elisa Bertino, Elena Ferrari, and Barbara Carminati for collaborating with us on the topics we discussed here. This work was partially supported by grants from the US Air Force Office of Scientific Research. We thank Robert Herklotz for his support.

## References

1. H. Chen, *Intelligence and Security Informatics and International Security: Information Sharing and Data Mining*, Springer, 2006.
2. T. Berners Lee, J. Hendler, and O. Lassila, "The Semantic Web," *Scientific Am.*, 17 May 2001.
3. M. Scannapieco et al., "Privacy-Preserving Schema and Data Matching," *Proc. ACM SIGMOD Int'l Conf. Management of Data*, ACM Press, 2007, pp. 653–664.
4. B. Thuraisingham, *Building Trustworthy Semantic Webs*, CRC Press, 2007.
5. T. Cadenhead, M. Kantarcioglu, and B. Thuraisingham, "Scalable and Efficient Reasoning for Enforcing Role-Based Access Control," *Proc. IFIP 11.3 Conf. Data and Applications Security*, IFIP, 2010.
6. R. Agrawal and R. Srikant, "Privacy-Preserving Data Mining," *Proc. ACM SIGMOD Int'l Conf. Management of Data*, ACM Press, 2000, pp. 439–450.
7. M. Kantarcioglu and J. Vaidya, "Secure Multiparty Computation Methods," *Encyclopedia of Database Systems*, 2009.

**Bhavani Thuraisingham** is the Louis A. Beecherl, Jr. I Distinguished Professor of Computer Science and director of the Cyber Security Research Center at the University of Texas at Dallas. Contact her at bhavani.thuraisingham@utdallas.edu.

**Latifur Khan** is an associate professor of computer science and director of the Data Mining/Database Laboratory at the University of Texas at Dallas. Contact him at latifur.khan@utdallas.edu.

**Murat Kantarcioglu** is an assistant professor and director of the Data Security and Privacy Laboratory at the University of Texas at Dallas. Contact him at muratk@utdallas.edu.

## Privacy-Preserved Social Network Integration and Analysis for Security Informatics

Christopher C. Yang, *Drexel University*  
Bhavani M. Thuraisingham, *University of Texas at Dallas*

Social network analysis (SNA) has been widely explored to support intelligence and law enforcement agencies in investigating the terrorist and criminal social networks. It is valuable in identifying terrorists, suspects,

subgroups, and their communication patterns.<sup>1</sup> Many related works on criminal and terrorist SNA have been published in the Intelligence and Security Informatics conference series ([www.isiconference.org](http://www.isiconference.org)). Although SNA has been proven to be important in security informatics, there are practical limitations in applying these techniques to conduct a large-scale analysis. Terrorist and criminal social network data is usually generated by intelligence and law enforcement agencies. Sharing across agencies is generally restricted, if not prohibited due to the privacy concerns. As a result, using such limited social network data diminishes the SNA performance. In some cases, results could incorrectly identify a criminal subgroup or be unable to identify a direct connection between terrorists. It is crucial to develop appropriate privacy-preserving social network algorithms that work with social network integration.

A number of researchers have published articles recently on preserving the privacy of social network data. However, not all these approaches apply to social network integration. In this article, we discuss the strength and limitations of the leading approach, which is based on anonymization models, and then introduce the subgraph generalization approach for social network integration.

## Anonymization Approach

Researchers developed anonymization models of preserving privacy for relational data about a decade ago. The objective is hiding the sensitive information such as personal identities but publishing the rest of the data—an anonymized version of relational data. However, a trivial linking attack can counter a naïve approach of simply removing names and identifications by using a set of

quasi-identifiers across multiple databases.<sup>2</sup> To protect against such attacks, the  $k$ -anonymity method ensures at least  $k$  records with respect to every set of quasi-identifier attributes are indistinguishable.<sup>2</sup> Other alternative methods such as  $l$ -diversity and  $m$ -invariance apply different constraints on anonymity.

Anonymization models developed for relational data cannot be directly applied on social network data, however, because social network data uses a graph representation rather than a tabular representation. In recent years, researchers have also developed anonymization models for preserving privacy on social network data such as  $k$ -candidate anonymity,  $k$ -degree anonymity, and  $k$ -anonymity.<sup>3</sup> These anonymization models of social networks remove the identities of nodes but retain and release the edges of the social network. However, simple active or passive attacks (known as neighborhood attacks) can infer the identities of nodes by solving a set of restricted isomorphism problems. For example, if adversaries have some knowledge about a target node's neighbors, they can reidentify the target node in the anonymized social network by extracting nodes that are isomorphic to the known target node.

By adding or deleting the social networks' nodes or edges (edge/node perturbation) in these anonymization models, adversaries can only have a confidence of  $1/k$  to discover the identities of nodes. We can analyze the global properties of the anonymized versions of the perturbed social networks. However, there are limitations of applying such privacy-preserving techniques on integrating and analyzing terrorist or criminal social networks. The perturbation techniques distort the social network's original structure to preserve privacy.

In some cases, the distortion is as high as 6 percent. Although the global properties might not vary substantially, some specific properties can be substantially changed. For example, the distance between two criminals can be reduced from  $n$  to two after perturbations. In addition, it is impossible to integrate these anonymized social networks because all identities are removed and there are no integration points between any pairs of anonymized social networks.

### **Subgraph Generalization Approach**

In light of such limitations, researchers have recently proposed the subgraph generalization approach for social network integration and analysis.<sup>4,5</sup> The identities of nodes in terrorist or criminal social networks are usually considered sensitive and will not be shared across different intelligence or law enforcement units. In practice, however, the identities of some terrorists or criminals are known across multiple units and even to the public. For example, Osama bin Laden is known to the public and multiple intelligence units. Depending on the privacy policy in individual units, some of these identities can be released. As a result, the nodes in social networks can be classified as sensitive nodes or insensitive nodes, which include the identities that can be shared across other units. The identities that must be preserved and not shared across units are considered sensitive nodes. The insensitive nodes will then be treated as the integration points in the social network integration process.

Given multiple social networks from different sources, the objective of integration and privacy preservation is incorporating the shared and

privacy-preserved data to achieve a higher SNA accuracy. Without integration, users can only conduct SNA on their own social network. The typical social network analyses are centrality measurement and clustering that involves measuring the distance between any two nodes.

Preparing social networks for integration involves two major steps. The first step is constructing subgraphs from a given social network where each subgraph must include one or more insensitive nodes as integration points. The second step is generating generalized information for each constructed subgraph, which is considered a generalized node.

### **Subgraph Construction**

To construct subgraphs from a social network, researchers have investigated a few techniques including  $K$ -nearest neighbor (KNN) and edge betweenness based (EBB) methods.<sup>4</sup> The KNN method takes an insensitive node as the centroid. A node  $v$  is assigned to one of the  $K$  subgraphs with the insensitive node  $v^c$  only if the shortest path between  $v$  and  $v^c$  is shorter than the shortest paths between  $v$  and other centroids. The EBB method removes edges with the highest betweenness iteratively to decompose a social network into a number of subgraphs with the constraint that each subgraph must have at least one insensitive node. Using either subgraph construction method, a number of subgraphs are generated from a social network and each subgraph is considered a generalized node of the generalized social network (see Figure 8). Two generalized nodes  $u'$  and  $v'$  are connected in the generalized social network if and only if a node in the corresponding subgraph of  $u'$  is connected to a node in the corresponding subgraph of  $v'$ .



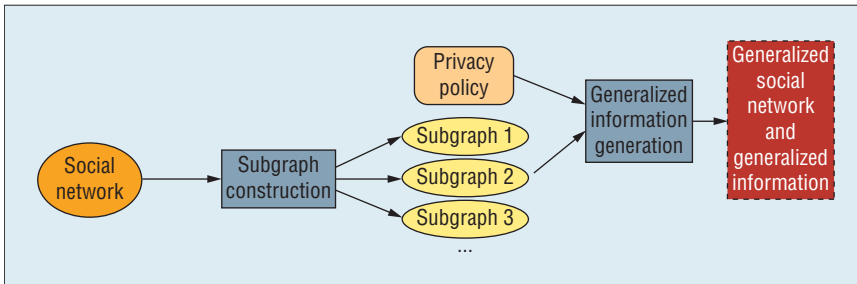


Figure 8. Creating a generalized social network and generalized information from a social network.

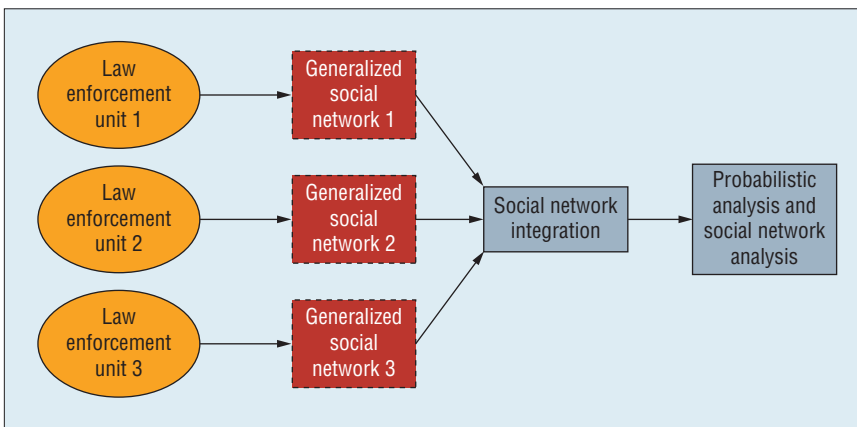


Figure 9. Social network integration and analysis.

## Generating Generalized Information

After constructing subgraphs, the sensitive information within a subgraph (generalized node) is not shared across intelligence or law enforcement units; only generalized information generated according to the privacy policy will be shared. The generated information describes the corresponding subgraph's general information that is useful for SNA. The more information we can share, the higher the information's utility, but this also poses a higher risk of a successful attack. The generalized information might include (but is not limited to) the number of nodes in a subgraph, the distribution of distances between all possible pairs of nodes in a subgraph, the number of shortest paths going through an insensitive node, the length of the shortest path between two insensitive nodes, the degree of insensitive

nodes, and the eccentricity of insensitive nodes. To determine what generalized information to share depends on the individual privacy policy and the SNA we want to conduct.

Given the generalized social networks and information from multiple intelligence and law enforcement units, we can conduct probabilistic analysis to integrate the shared information with the owned social network (see Figure 9). By integrating the generalized subgraphs' probabilistic models, we compute the probability of distance between any two nodes to estimate the centrality measures. Experiments have shown that the subgraph generalization approach can reduce the errors in estimating closeness centrality from 40 to 20 percent.<sup>4,5</sup>

Research has shown that SNA is valuable in knowledge discovery of terrorist and criminal interaction patterns and identifying suspects.

Unfortunately, its effectiveness always depends on the completeness of each intelligence or law enforcement unit's social network. Efforts on the proposed subgraph generalization approach are ongoing and promising in integrating social networks and preserving privacy. ■

## References

1. H. Chen and C.C. Yang, *Intelligence and Security Informatics: Techniques and Applications*, Springer Verlag, 2008.
2. P. Samarati, "Protecting Respondents' Identities in Microdata Release," *IEEE Trans. Knowledge and Data Eng.*, vol. 13, 2001, pp. 1010–1027.
3. K. Liu and E. Terzi, "Towards Identity Anonymization on Graphs," *Proc. ACM SIGMOD*, ACM Press, 2008.
4. X. Tang and C.C. Yang, "Generalizing Terrorist Social Networks with K-Nearest Neighbor and Edge Betweenness for Social Network Integration and Privacy Preservation," *Proc. IEEE Int'l Conf. Intelligence and Security Informatics*, 2010.
5. C.C. Yang, X. Tang, and B.M. Thuraisingham, "Social Networks Integration and Privacy Preservation using Subgraph Generalization," *Proc. AMC SIGKDD Workshop CyberSecurity and Intelligence Informatics*, 2009.

**Christopher C. Yang** is an associate professor in the College of Information Science and Technology at Drexel University. Contact him at [chris.yang@drexel.edu](mailto:chris.yang@drexel.edu).

**Bhavani Thuraisingham** is the Louis A. Beecherl, Jr. I Distinguished Professor and director of the Cyber Security Research Center at the University of Texas at Dallas. Contact her at [bhavani.thuraisingham@utdallas.edu](mailto:bhavani.thuraisingham@utdallas.edu).

**cn** Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.