Synergistic and Perceptual Intrusion Detection with Reinforcement (SPIDER)

Patrick Miller*
Eastern Washington University
Cheney, WA 99004

John Mill†
Spokane Falls Community College
Spokane, WA 99224

Atsushi Inoue‡
Eastern Washington University
Cheney, WA 99004

Abstract

An intrusion detection system based on multi-agent architecture and soft computing is introduced. This consists of multiple, autonomous agents and a single blackboard mechanism in order to facilitate intrusion detection for a target computer network. Within this framework, each agent performs its own tasks asynchronously. The blackboard mechanism consistently manages access to the constituent agents. Each agent thus can be implemented arbitrarily as long as the interface to the blackboard is supported. Currently, Fuzzy Logic, Neural Networks and Support Vector Machines are being studied as underlying methodologies for the autonomous agents. In addition, a single agent can be capable of reinforcement learning with the reinforcement signal provided by the blackboard. The system development is being pursued using both software and hardware implementations.

Introduction

Intrusion Detection, Network Security and Artificial Intelligence

As the Internet becomes increasingly dominant, securing computer systems from unwanted intrusion is becoming a top priority within many IT departments. Statistics show that the number of incidents of computer network security compromise has increased by a factor of over 300 between the years 1990 and 2002 (CERT 2003). Table 1 shows the type of intrusions, their descriptions and effects commonly identified.

Table 2 shows the type of solutions to prevent or deal with intrusions that are commonly taken into account (Axelsson 1999a). The latter primarily require recording and tracing through massive quantities of activity logs. If the study of the activity logs shows signs of intrusion, certain actions are taken to prevent or to eliminate the intrusion. This is not a trivial task due to the massive quantity of data and the requirement for the system to respond in real-time.

Currently, system administrators must rely on their experience to configure the intrusion management utilities described above. Typically, they configure the utilities so that monitoring processes are placed on critical sites (e.g. fire walls, gateways and web servers) and centralized at one site (their primary workstation). To realize this, the utilities must have a high degree of inter-operability and the capability to provide summaries. Further, security assurances on these utilities themselves need to be provided in one way or the other. Despite the advantages provided by these tools, system administrators are under increasing pressure due to the massive quantity of network traffic, dynamic network configurations and increased intrusion attempts.

In this context, some real-time intelligent processing techniques acting upon network traffic information are expected to provide assistance to the system administrator in his/her decision making process. The task of intrusion detection is performed with certainty factors. Justifications of determining such certainty factors vary depending on sources of inputs, expertise and perspectives. Usually, (Bayesian) probability is incorporated within such an intelligent processing in one way or the other as the underlying representation framework.

In making (Bayesian) probabilistic reasoning, general and broad-based information about the population characteristics (i.e. prior probability of an event occurring) needs to be taken into account. When the prior is neglected, the conditional probability of the event occurring given a certain observation (i.e. a rule for Bayesian inference) results in either insignificant or inaccurate results regardless of a considerable size of samples (i.e. base rate fallacy, a.k.a. representative heuristics (Kahneman & Tversky 1973)). Some studies (Gigerenzer & Hoffrage 1995) showed that subjective estimates based on frequencies clearly overcome normative predictions and they often come close to a hit rate of 80% or better.

*E-mail: patrick@doriathproject.com
†E-mail: johnmi@spokanefalls.edu
‡ Director of Inland Northwest Security System Initiative (INSSI) within Department of Computer Science at Eastern Washington University. E-mail: atsushi.inoue@ewu.edu

Type	Description	Effect
Denial of services (DoS)	Disables or floods network devices or infrastructure (E.g. ping of death, SQL attacks)	Slowing down or disabling network services
Probing	Sniffing the network gateways and network services	No harmful effects Precursor to malignant attacks
Brute force intrusion	Unauthorized system access, e.g., illegal privilege login	Loss or alternation of information

Table 1: Classification of Network Intrusions

Solution	Example
Prevention	Fire walls Proxy servers Virtual private networks (VPN) Cryptography
Monitoring	Network packet monitoring: TCPdump (LBNL 2003) Traffic analysis and visualization: CoralReef (CAIDA 2003) Intrusion detection systems: Snort (Caswell & Roesch 2003)

Table 2: Examples of Security Mechanisms

Obviously, it is very difficult to obtain priors for use in intrusion detection(Axelsson 1999b). Yet, competent system administrators are capable of successfully detecting intrusions. On the other hand, a widely used intrusion detection system, Snort, utilizes a pattern matching scheme based on regular expression. A survey (Axelsson 1999a) indicates that, in general, knowledge-based approaches are among the most often used for intrusion detection.

Current research has focused on a model for team decision making using Distributed Artificial Intelligence (Kang, Waisel, & Wallace 1998). This indicates that team decision making performs better than individual decision making, especially in complex tasks (though there is an expense associated with resolving conflicts among team members).

Research Statement

The primary objective of Synergistic and Perceptual Intrusion DEtection with Reinforcement (SPIDER) is to seek a more efficient and effective intrusion detection framework by promoting synergistic effects of autonomous software agents and system administrators. The efficiency and effectiveness are considered with respect to *inter-operability*, *information summarization* and *self-security*. The agents are expected to be knowledge-based and, preferably, adaptive (i.e. capable of reinforcement learning). They are also expected to be able to handle uncertainty. A simple team decision making model with a conflict resolution scheme is studied. It needs to be compliant with various types of agents and human beings, i.e. system administrators, through a certain user interface.

A secondary objective of SPIDER is a significant contribution to our new Cyber Security program (see Appendix) including, but not limited to: course material, promotion of collaboration between faculty and industry, dissemination such as tutorials and workshops for the local chapter of the ACM and regional special interest groups (SIG), and master's thesis projects. It serves as a testbed for studying techniques of Artificial Intelligence (e.g., reasoning, learning, multi-agents) and network security.

Multi-Agent, Blackboard-Based System Architecture

Figure 1 shows an overview of SPIDER whose architecture is a multi-agent, blackboard-based system (Russell & Norvig 2003). Its subsystems include:

1. *Blackboard subsystem* (BB) - It maintains information regarding intrusion detection in the network. It allows agents to access and modify information asynchronously. Support logic (Baldwin, Martin, & Pilsworth 1995) is used for representation and inference about intrusions.
2. *Blackboard management subsystem* (BBM) - It has two tasks: the combination of decisions (i.e. intrusion detection) made among multiple agents and housekeeping such as elimination and compression of outdated information.
3. *User interface* (UI) - This provides a graphic user

interface (GUI) in order for a user, i.e. a system administrator, to make a decision whether or not a certain action to handle an intrusion is necessary. It provides visualization of network traffic, processes and other significant information stored on BB. This subsystem also controls the modules performing actions in order to handle intrusions (e.g., shutting off machines, filtering a certain type of network packet, etc.).

4. *Intrusion detection agent* (IDA) - A single agent autonomously and independently performs its intrusion detection task. IDAs which belong to SPIDER will have their results reflected on BB.

An IDA detects intrusions in real-time based on patterns in histograms that are generated from input streams such as network packets, datagrams, E-mail streams and OS process tables. Histograms are generated by taking frequencies of sequentially occurring events sensed from input streams within a certain period. Histograms are taken as the canonical form for input streams and sampling processes of the IDA. The following classification can be made in terms of the type of input streams:

1. Host-based IDA - Generates histograms from OS process tables.
2. Network-based IDA - Generates histograms from network packet streams.
3. Application-based IDA - Generates histograms from datagrams and/or E-mail streams(Inoue & Ralescu 1999).

Figure 2 shows a snapshot of SPIDER with a variety of input streams. The system administrator monitors various types of intrusions in an integrated manner on his/her workstation through UI. Computational models of IDAs include fuzzy logic, neural networks, probabilistic reasoning, support logic and string matching based on regular expressions.

Decision Making Model in SPIDER

Within SPIDER, the system administrator and agents collaboratively make decisions as to whether or not intrusions are detected. All decisions made by those agents and the system administrators are recorded on BB. BBM then aggregates them in order to determine the most appropriate actions. In the following, the decision model underlying the BBM is briefly described.

The main idea is to make a decision by a team consisting of the IDAs and the system administrator. When considering the efficient frameworks of decisions (reasoning) made by IDAs, e.g., probability, Dempster-Shafer and fuzzy sets, as well as the histograms generated in real-time, an underlying representation framework is necessary. In addition, decisions made by human beings (i.e. system administrators) are involved. Consequently, the team decision needs to be perceptual due to the involvement of human beings with the various reasoning frameworks of the IDAs.

Consider a set of decisions D corresponding all intrusions and normal cases (i.e. no intrusions). Without loss of generality, we assume that at least one of the IDAs or system administrators is capable of making decisions defined within D. This assumption prevents decisions that are never made. Then we consider the following decisions made by individual decision makers:

1. *Crisp decision*: a particular decision $x_c \in D$ is made (i.e. probability $P(x_c) = 1$).
2. *Probabilistic decision*: a decision $x_p \in D$ associated with a point probability $P(x_p)$ such that $P(x_p) + P(\overline{x_p}) = 1$ is made.
3. *Support decision*: a decision $x_s \in D$ associated with a support pair, i.e. an interval of probability, $(P_l(x_s), P_u(x_s))$ where $P_l(x_s) \leq P_u(x_s)$ is made. Consequently, the support pair for the complement $(P_l(\overline{x_s}), P_u(\overline{x_s}))$ is determined such that $(1 - P_u(x_s), 1 - P_l(x_s)) = (P_l(\overline{x_s}), P_u(\overline{x_s}))$.
4. *Fuzzy decision*: a decision $x_f \in D$ associated with a fuzzy probability $P_f(x_f)$, a fuzzy set defined over $[0, 1]$, is made. (E.g., $P_f(x_f)$ = 'high' where 'high' is a fuzzy set defined over $[0, 1]$).

Next, consider how such various types of decisions made by IDAs and system administrators (decision makers) are combined. Without loss of generality, assume that n decision makers making various types of decisions belong to SPIDER. Then the team decision x associated with a probability $P(x)$ is obtained by aggregating $P_i(x)$ where $i = 1 \ldots n$ (i.e. the probability that the i-th decision maker makes decision x) such that

$$P(x) = \mathcal{H}_w(P_1(x), \ldots, P_n(x)) \quad (1)$$

where w is the vote weight (usually normalized, i.e. $\sum_i w_i = 1$) (representing the influence on the decision making). $\mathcal{H}$ is a notation indicating a generic aggregation combination operation such that $\mathcal{H} : [0, 1]^n \mapsto [0, 1]$.

Point Combination

Many options for the combination operation can be considered. Currently, a team decision making model with a simple weight averaging combination such that

$$P(x) = \sum_{i=1}^{n} w_i \cdot P_i(x) \quad (2)$$

where $\sum_i w_i = 1$ (i.e. the summation of normalized votes) is currently studied by Miller (Miller & Inoue 2003). In this study, IDAs make only crisp or probabilistic decisions.

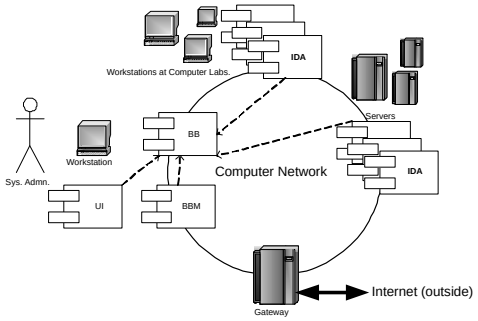

Figure 1: Overview of SPIDER

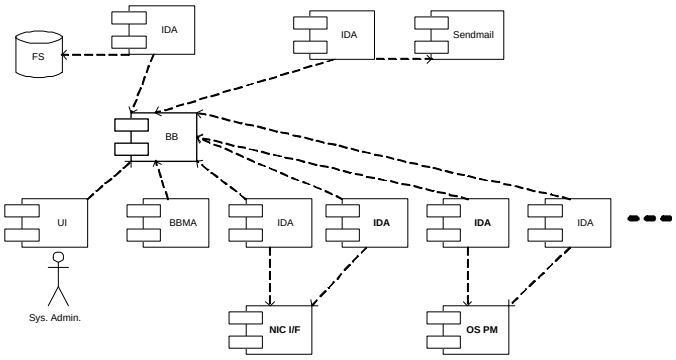

Figure 2: Snapshot of Synergistic Intrusion Detection

How can support decisions and fuzzy decisions be handled? The issue arises in connection with using Equation 2 in this case. A simple solution is to take a middle point

$$\hat{P}_i(x) = \frac{P_l^i(x) + P_u^i(x)}{2}$$

for the support decisions, and the deffuzification

$$\hat{P}_i(x) = \frac{\sum_{y \in [0,1]} \mu_{P_i}(y) \cdot y}{\sum_y \mu_{P_i}(y)}$$

for the fuzzy decisions to let point probabilities represent them respectively.

Support Combination

Alternatively, we are currently studying a more generic combination operation utilizing Mass Assignment Theory (MAT) (Baldwin, Martin, & Pilsworth 1995). MAT provides a framework for managing the correspondence between fuzzy sets and probability distributions using mass assignment as a mediator. Here, we consider the combination of decisions utilizing MAT similar to the construction of computational perception (Inoue & Ralescu 2000).

(Definition 1) Let S be a sample space. Then a *mass assignment* (MA) m_S associated with S is a function from the power set $\mathcal{P}(S)$ to an interval of real numbers such that

$$m_S : \mathcal{P}(S) \mapsto [0,1]$$

and

$$\sum_{A \subseteq S} m_S(A) = 1$$

(Definition 2) $A \subseteq S$ is called a *focal element* for mass assignment m_S if

$$m_S(A) > 0$$

MAT provides the following correspondences among probability distributions, mass assignment, and fuzzy sets:

1. *MA and probability*:

$$P_S(x) = \sum_{A \subseteq S, x \in A} P_A(x) \cdot m_S(A) \qquad (3)$$

where P_S is a probability distribution on S, m_S is a mass assignment over S and P_A is a probability distribution on A (often called a selection rule). The selection rule represents bias (i.e. preference) on elements within A. The selection rule without any bias is $P_A(x) = \frac{1}{|A|}$ (the least prejudged distribution).

2. *MA and fuzzy sets*: Let $F = x_1/\mu_1 + \cdots + x_n/\mu_n$ be a fuzzy subset over S. We denote $\mu_i = \mu_F(x_i)$ and without loss of generality we assume

$$1 = \mu_1 \geq \cdots \geq \mu_n \geq \mu_{n+1} = 0$$

Then a MA with nested focal elements $\{x_1, \ldots, x_i\}$ for $i = 1, \ldots, n$ can be derived as

$$m_S(A) = \begin{cases} \mu_i - \mu_{i+1} & \text{if } A = \{x_1, \ldots, x_i\} \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

3. *Probability and fuzzy sets*: The mapping between fuzzy sets and probability distributions via a MA is obtained from above two such that

$$P_S(x_k) = \sum_{i=k}^{n} P_A(x_k) \cdot (\mu_i - \mu_{i+1}) \quad (5)$$

Let D_j be a set of decisions that the j-th decision maker makes. Note that $D_j \subseteq D$ and $\bigcup_{j \in \{1 \ldots n\}} D_j = D$ (D is a set of decisions made by a team within SPIDER). Suppose that the j-th decision maker has a normalized vote w_j (i.e. $\sum_{j \in \{1 \ldots n\}} w_j = 1$). Then the probability that a decision $x \in D$ is made by a team within SPIDER is given by

$$P_D(x) = \sum_{D_j \subseteq D, x \in D_j} P_{D_j}(x) \cdot w_j \quad (6)$$

This is derived from Equation 3 by treating focal element $A = D_j$ and MA $m_D(A) = m_D(D_j) = w_j$. $P_{D_j}(x)$ is a probability associated with a decision made by the j-th decision maker. As mentioned above, this probability can be either one of the following:

1. *Crisp Decision*: $P_{D_j}(x) = 1$ and $P_{D_j}(\overline{x}) = 0$
2. *Probabilistic Decision*: Let the probability for a decision x be $P_{D_j}(x) = p$. Then $P_{D_j}(\overline{x}) = \frac{1-p}{|D_j|-1}$ (assuming no bias).
3. *Support Decision*: Consider the following procedure:
   (a) Generate MA m_j corresponding to a support pair (p_l, p_u) for a decision x such that

$$m_j(\{x\}) = p_l$$
$$m_j(\{\overline{x}\} = D_j - \{x\}) = 1 - p_u$$
$$m_j(D_j) = p_u - p_l$$

   (b) Compute $P_{D_j}(x)$ from m_j using Equation 3. Assume the least prejudged distribution $\frac{1}{|A|}$, where A is a focal element of m_j, for the selection rules unless otherwise specified.
4. *Fuzzy Decision*: Let F_p be a fuzzy probability such that $P_{D_j}(x) = F_p$. By the representation (decomposition) theorem (Klir & Yuan 1995), we obtain

$$F_p = \bigcup_{\alpha \in [0,1]} \alpha \cdot F_\alpha$$

where $\alpha \cdot F_\alpha$ is a special fuzzy set whose membership function is given by $\mu_{\alpha \cdot F_\alpha}(x) = \alpha$ for all $x \in F_\alpha$, F_α is a crisp set (α-cut)consisting of elements x such that $\mu_{F_p}(x) \geq \alpha$. Notice that any α-cut F_α is an interval within $[0, 1]$ provided that fuzzy set F_p is convex.

MA m_{F_p} corresponding to F_p can be obtained such that

$$m_{F_p}(F_{\alpha_i}) = \alpha_i - \alpha_{i+1}$$

where, without loss of generality, α is sorted in non-increasing order such that

$$1 = \alpha_1 \geq \cdots \geq \alpha_n \geq \alpha_{n+1} = 0$$

and F_α are the only focal elements for m_{F_p}. This leads to a collection of possible support decisions such that

$$(p_l^i = MIN[F_{\alpha_i}], p_u^i = MAX[F_{\alpha_i}])$$

Lastly, m_j for the collection of possible support decisions from F_p is obtained such that

$$m_j(\{x\}) = \sum_{i=1}^{n} p_l^i \cdot m_{F_p}(F_{\alpha_i})$$
$$m_j(\{\overline{x}\} = D_j - \{x\}) = \sum_{i=1}^{n} (1 - p_u^i) \cdot m_{F_p}(F_{\alpha_i})$$
$$m_j(D_j) = \sum_{i=1}^{n} (p_u^i - p_l^i) \cdot m_{F_p}(F_{\alpha_i})$$

Then $PD_j(x)$ is computed from m_j by using Equation 3.

Reinforcement of Adaptive IDAs

Reinforcement of SPIDER is performed in a manner of propagating rewards or penalties from BB to all IDAs with adaptive capabilities. This procedure is outlined below:

1. A system administrator performs a certain action for decision $x \in D$.
2. For decision x, determine its reward $r(x) = \gamma \cdot (1 - P_D(x))$ where $\gamma \in [0, 1]$ is a constant which determines a learning rate.
3. For other decisions x_i, where $x_i \neq x$, determine its penalty $r(x_i) = -r(x) \cdot \frac{P_D(x_i)}{\sum_{i, x_i \neq x} P_D(x_i)}$
4. Update the probability $P_D(x_i) \rightarrow P'_D(x_i)$ for all $i \in \{1 \cdots n\}$ by $P'_D(x) = r(x_i) + P_D(x_i)$.
5. Obtain the corresponding updated (normalized) vote $w_i \rightarrow w'_i$ by following Equation 6.
6. Determine rewards or penalties $r_i^w = w'_i - w_i$ for all decision makers.
7. Propagate r_i^w to all adaptive IDAs.
8. Adjust the vote w_i if the i-th IDA has adapted itself to improve performance (this is applied only when a penalty is applied).

IDAs Currently Studied

As mentioned above, IDAs can take any computational model as long as it is compliant with the scheme mentioned above. Currently, the following computational models are being studied:

1. Support Logic and Mass Assignment Theory - This provides an underlying computational framework of Bayesian belief networks, support logic, Dempster-Shafer theory of evidence, fuzzy logic and neural networks (Inoue 2003). Fuzzy Relational Inference Language (FRIL) provides a logic programming environment with S-expression notations. The goal of this research is to implement IDAs as embedded systems, i.e. hardware interface boards.
2. Support Vector Machines and an extension of support vectors by using fuzzy sets - This framework provides an instance-based classification approach that suits many problem domains such as text classification (Mill 2002) and intrusion detection where knowledge acquisition is very difficult. To handle 'near errors', an extension of support vectors with fuzzy sets is currently under study (Mill & Inoue 2003).
3. Self-Organizing Map (SOM) - Miller developed a simple unauthorized privilege access (root password) detection using SOM (a single IDA). This was extended to the network packet intrusion detection problem together with the incorporation of SPIDER architecture (Miller & Inoue 2003).
4. Other candidate computational models - Genetic algorithms, artificial life, and artificial immune systems are currently being studied.

Self-Security

Self-security becomes a big issue for any distributed system. The fundamental advantage for systems such as SPIDER is that it is hard for an intrusion to penetrate the entire system due to its distributed nature. The major potential vulnerabilities are DoS attacks (as the network is flooded, SPIDER cannot exchange information among IDAs). The solution to this problem is to set up SPIDER on a totally isolated network. This is feasible considering the bandwidth that SPIDER uses (very small). It is simple to set up such an isolated network either physically installing extra network interface cards or virtually configuring a separate, isolated network (i.e. virtual private network (VPN)).

Conclusion

A distributed intrusion detection scheme, SPIDER, is introduced. It has a multi-agent, blackboard architecture which enables integrated intrusion detection at all levels. A synergistic effect is expected by having a variety of agents for each intrusion. A simple decision making model using fuzzy sets and voting is introduced in order to combine intrusion detection results generated by IDAs. A simple distributed reinforcement scheme is outlined.

SPIDER is expected to impact on our cyber security program and to serve as a testbed for multiple disciplines such as Artificial Intelligence and Network Security. This provides many opportunities for course and master's thesis projects. This will also promote collaboration between universities and industry as well as supporting community outreach.

To reflect the current trends in intrusion detection and management, significant future work involves the extension of SPIDER to be able to manage intrusions by taking actions. The key issue is the configuration of actions reflecting the decisions with uncertainties, i.e. reflection of P_D to the parameters of actions, e.g., the blend control framework of unmanned helicopter (Sugeno *et al.* 1995).

Acknowledgment

This work is in part supported by the Congressional Appropriation of Technology Initiative for New Economy (TINE) for the development of a new School of Computing and Engineering Science and NSF planning grants (NSF 0230590) for the development of a new Software Engineering Technology program.

Appendix: Cyber Security Program

The primary missions of the Cyber Security program at EWU are to train a cyber security workforce and to serve as the leading think-tank for the Inland Northwest region. The program development consists of the following:

1. New course and academic module development in three tracks
   (a) Theory and foundation track – Cryptography and cryptographic protocols
   (b) Security engineering track – network security, OS security, web security, and information warfare
   (c) Application track – intrusion detection within artificial intelligence, secured programming in software engineering and programming core.
2. Enhancement of degree programs - Revision of the BS and the MS degrees in Computer Science to include the cyber security within the core. A new undergraduate degree program, Software Engineering Technology (Rodriguez-Marek *et al.* 2003) will be offered in Fall of 2003 with an emphasis on 'hands-on' network security and embedded systems development.
3. New applied research programs - Synergistic collaborations among faculty members, students, government research laboratories and local industry.

4. Community outreach - The Inland Northwest Security System Initiative (INSSI) was established in order to promote applied research and collaborative efforts (e.g. workshops and tutorials) for cyber security among universities, industries and research laboratories in this region.

References

Axelsson, S. 1999a. Intrusion detection systems: A taxonomy and survey. Technical Report 99-15, Dept. of Computer Engineering, Chalmers University of Technology, Sweden.

Axelsson, S. 1999b. On a difficulty of intrusion detection. In *Procceedings of 2nd Intl. Workshop on Recent Advances in Intrusion Detection (RAID'99).*

Baldwin, J. F.; Martin, T.; and Pilsworth, B. 1995. *FRIL: Fuzzy and Evidential Reasoning in AI.* Research Studied Press.

CAIDA. 2003. CoralReef, http://www.caida.org/tools/measurement/coralreef/.

Caswell, B., and Roesch, M. 2003. http://www.snort.org/.

CERT. 2003. http://www.cert.org/.

Gigerenzer, G., and Hoffrage, U. 1995. How to improve bayesian reasoning without instruction: Frequency formats. *Psychological Review* 102(4):684–704.

Inoue, A., and Ralescu, A. L. 1999. E-mail classification reflecting user perceptions. In *Computational Intelligence and Learning Workshop on Computational Intelligence for User Modeling.*

Inoue, A., and Ralescu, A. L. 2000. Construction of computational perception using mass assignment theory. In *International Conference on Information Processing and Management of Uncertainty in Knowledge-based Systems*, 1427–1434.

Inoue, A. 2003. Perceptual intrusion detection system. In *North American Fuzzy Information Processing Society.* submitted.

Kahneman, D., and Tversky, A. 1973. On the psychology of prediction. *Psychological Review* 80:237–251.

Kang, M.; Waisel, L. B.; and Wallace, W. A. 1998. Teamsoar: A model for team decision making. In Prietula, M. J.; Carley, K. M.; and Gasser, L., eds., *Simulating Organizations: Computational Models of Institutions and Groups.* MIT Press. 23–45.

Klir, G. J., and Yuan, B. 1995. *Fuzzy Sets and Fuzzy Logic: Theory and Applications.* Prentice Hall.

LBNL. 2003. http://www.tcpdump.org/.

Mill, J., and Inoue, A. 2003. An application of fuzzy support vectors. In *North American Fuzzy Information Processing Society.* submitted.

Mill, J. 2002. Support vector machines, n-gram kernels, and text classification. Master's thesis, Eastern Washington University.

Miller, P., and Inoue, A. 2003. Collaborative intrusion detection system. In *North American Fuzzy Information Processing Society.* submitted.

Mitchell, T. M. 1997. *Machine Learning.* McGraw-Hill.

Rodriguez-Marek, E.; Brzoska, M. A.; Koh, M.; Loendorf, W.; and Inoue, A. 2003. Developing a software engineering technology program. In *American Society for Engineering Education Annual Meeting.* accepted.

Russell, S., and Norvig, P. 2003. *Artificial Intelligence: Modern Approach.* Prentice Hall, 2nd edition.

Sugeno, M.; Hirano, I.; Nakamura, S.; and Kotsu, S. 1995. Development of an intelligent unmanned helicopter. In *IEEE International Conference on Fuzzy Systems*, volume 5, 15–16.