

Multicriteria Security System Performance Assessment Using Fuzzy Logic

William L. McGill, PE, CRE

College of Information Sciences and Technology
The Pennsylvania State University, University Park, PA 16802
(814) 308-3854
wmcgill@ist.psu.edu

Bilal M. Ayyub, PhD, PE

Center for Technology and Systems Management
Department of Civil and Environmental Engineering
University of Maryland, College Park, MD 20740
(301) 405-1956
ba@umd.edu

Modern security problems focus on sensibly allocating resources to decrease the magnitude of potential hazards, decrease the chances of adversary success given an attempt, or minimize loss following a successful attack. The focus of this paper is on developing a simple, yet analytically sound tool that facilitates rapid assessments of security system non-performance in terms of probability of adversary success at the facility or asset level using concepts from fuzzy logic. Beginning with a short overview of how security system performance fits within an overall security risk analysis frame-work, this paper presents the basic concepts of fuzzy systems and applies them to develop a model that approximates the true relationship between defensive capabilities and probability of adversary success. A simple example demonstrating the proposed model to support decision making accompanies this discussion. This paper concludes with a strategy for implementation of the proposed model in an operational setting.

Keywords: risk analysis, fuzzy systems, fuzzy logic, probability of adversary success, homeland security

1. Introduction

Modern security problems focus on sensibly allocating resources to decrease the magnitude of potential adverse events (e.g., malicious attack), decrease the chances of perpetrator success given an attempt, or minimize loss following a successful incident [1,2]. The modern risk environment challenges security and risk professionals to allocate resources among a wide array of countermeasures and consequence mitigation strategies so as to manage a decision-maker's risk exposure in an affordable, cost-effective manner. Accordingly, a sound business case for security investments must be accompanied by meaningful measures for the return on investment, such as risk-based benefit-to-cost ratios [3]. In this context, the knowledge of a decision-maker's risk exposure before and after implementation of a risk mitigation strategy is essential for measuring the benefit of a proposed investment.

Risk is a multidimensional concept. Though commonly expressed numerically as the product of the probability of occurrence and expected consequences associated with an adverse event, risk is generally defined as the set of all triplets containing an initiating event, the probability of that initiating event occurring, and its associated consequences [4]. In the context of risk analysis for critical infrastructure protection, risk is often expressed as the Cartesian product:

$$Risk = Threat \times Vulnerability \times Consequence \quad (1)$$

where *threat* describes a set of adverse initiating events (e.g., explosive attack), *consequence* describes the spectrum of losses that can be felt by the victims following their occurrence, and *vulnerability* comprises a set of system or target weaknesses (e.g., security, hardness) that can be exploited by an adversary to achieve a given degree of loss or harm. While Eq. 1 is broadly applicable to natural hazards, the human aspect of security hazards afford decision makers the added option to influence the probability of attack and probability of success given attack (e.g., security vulnerability) via deterrence and security countermeasures. That is, security risks can be reduced by improving the performance of a security system.

The focus of this paper is on developing a simple, yet analytically sound model to assess the probability of adversary success at a facility or asset with respect to a given initiating event using concepts from fuzzy logic. Beginning with a short overview of how security system performance fits within the overall security risk analysis paradigm, this paper presents the basic concepts of fuzzy systems and applies them to approximate the true functional relationship between the effectiveness of six security system capabilities and probability of adversary success. A simple example demonstrating the proposed model to support decision making accompanies this discussion. This paper concludes with a strategy for implementation of the proposed model in an operational setting.

2. Analytical Context

In general, the risks arising from security events (e.g., malicious attacks) can be expressed quantitatively in terms of the annual rate of exceeding a given level of loss c , $\lambda(C > c)$ as:

$$\lambda(C > c) = \lambda_A \sum_i \Pr(C > c | S, A_i) \Pr(S | A_i) \Pr(A_i) \quad (2)$$

where λ_A is the annual rate of attack occurrence in general (in number of events per year), $\Pr(A_i)$ is the probability of specific initiating event A_i given that an attack has occurred, $\Pr(S | A_i)$ is the probability of adversary success (S) given the occurrence of initiating event A_i , $\Pr(c > C | S, A_i)$ is the probability of realizing loss C that exceeds c given adversary success, and the summation is taken over an exhaustive and disjoint set of initiating events $A_i \in A$ (see for example [1,2]). In the context of security risk analysis, the desired measure of security system non-performance is the probability of adversary success given attempt, $\Pr(S | A_i)$. Other sources have described security system performance in terms of the complementary event $\Pr(\bar{S} | A_i)$, which has been referred to in the literature as security system effectiveness [5,6,7].

The performance of a security system under load depends on the nature of the specific initiating events considered. Initiating events can be specified generically in terms of an attack type (e.g., explosive attack), with greater resolution in terms of a pairing of attack type with a target (e.g., explosive attack against chemical tank), or more specifically in terms of a specific attack profile (e.g., truck bomb against chemical tank via rear access road). In all cases, the set of plausible initiating events must be exhaustive, and the choice of analytical resolution should only be as specific as needed to support decision making [8].

According to Manunta, the state of security is a complex function of the characteristics of the target, its methods of defense, and adversary intentions and capabilities [9]. As an engineered system, in principle the performance of a security system under stress can be determined through detailed systems modeling and analysis techniques. For example, a complex security systems model would yield insights into how different hardware, software, and human components interact to defeat an adversary, which would then support component and system level resource allocation and design decisions. Unfortunately, this level of resolution requires a significant investment in analytical resources to construct and validate complex security system models for each initiating event of interest.

Fortunately, many security decision problems do not require such high-resolution analysis; rather, only high level assessment of security system performance is needed to quantify overall security risk. For example, local governments may want to leverage simple analysis tools in order to obtain security-related information for many facilities that supports local or regional risk assessment. Perhaps this would explain the popularity of the CARVER methodology for vulnerability assessment within the security community [10]. The CARVER method identifies six vulnerability factors (i.e., Criticality, Accessibility, Recuperability, Vulnerability, Effect, and Recognizability), and subjectively assigns a value on the scale of 0 to 10 to each. The overall CARVER score is then taken as the sum of the scores assigned to the six criteria.

While the CARVER method arguably provides a quick and dirty means for ranking potential targets based on vulnerability, its additive and inherently non-probabilistic nature does not produce results that can support security risk assessment as required by Eq. 2. What is needed is a method with the front-end appeal of the CARVER methodology yet is built atop an analytically sound back-end that can relate subjective assessments of security system capabilities to information that can support quantitative risk analysis. To the authors' knowledge, no such tool currently exists. The remainder of this paper is devoted to developing such a model based on the principle of fuzzy systems and fuzzy logic. The result is a security performance assessment tool whose back-end leverages established mathematical techniques to approximate the probability of

adversary success according to the subjective assessment of multiple security effectiveness criteria.

3. Background on Fuzzy Systems

3.1. Fuzzy Numbers and their Membership Functions

The basic building blocks of fuzzy logic are linguistic variables described by fuzzy numbers. A linguistic variable is one that takes on linguistic values such as “strong” for a variable describing impact resistance of a crash barrier and “secure” for a variable defining the state of protection [11]. That is, a linguistic variable takes on values that have clear intension (“the barrier is strong”), but with a vague extension (we do not know the true failure load). In contrast to a crisp number whose value is precisely defined, a fuzzy number is a fuzzy set defined on the set of real numbers whose numeric meaning is vaguely defined [12]. For example, the word “likely” as a statement about probability and the word “catastrophic” as a statement about potential consequences are both fuzzy numbers in the sense that they express magnitude without precise quantification.

The degree of belonging or membership of a certain numeric value x to a fuzzy number X is characterized by a membership function $m(x)$ that assigns a value on the domain $[0,1]$, where a membership of 1 indicates that x fully belongs to X , a membership value of 0 indicates that x does not belong to X , and values in between indicate that x partially belongs to X . Figure 1, for example, shows a series of fuzzy numbers representing various degree of probability as derived from the data of Lichtenstein and Newman [13] assuming a single-valued core positioned at the median and the support spanning the range of responses for each probability phrase.

The general form of a membership function for a fuzzy number X is as follows:

$$m(x) = \begin{cases} m_L(x) & a < x < b \\ 1 & b \leq x \leq c \\ m_R(x) & c < x < d \\ 0 & \text{Otherwise} \end{cases} \quad (3)$$

where $m_L(x)$ and $m_R(x)$ are, respectively, non-decreasing and non-increasing functions of x on the domain $[0,1]$. The fuzzy number described by the membership function in Eq. 3 is known as a *generalized left-right fuzzy number*, or GLRFN [14]. The *core* of a fuzzy number X is defined as the interval where $m(x) = 1$ (i.e., $[b, c]$ in Eq. 3), and consists of all values x that definitely belong to X . The *support* of a fuzzy number is defined as the interval where $m(x) > 0$ (i.e., (a, d) in Eq. 3), and consists of all values x that have at least a partial belonging to X .

In practice, fuzzy numbers are commonly represented in a simplified or approximate form known as a *trapezoidal fuzzy number*, or TrFN. As with a GLRFN, a TrFN has four distinct points (a,b,c,d) , where the interval $[b,c]$ defines the core and the interval $[a,d]$ defines the support. The membership function of a TrFN is *linearly increasing* on the interval $[a,b]$, *constant* on the interval $[b,c]$, and *linearly decreasing* on the interval $[c,d]$. That is:

$$m_L(x) = \left(\frac{x-a}{b-a} \right) \quad (4)$$

$$m_R(x) = \left(\frac{d-x}{d-c} \right) \quad (5)$$

It is common practice to denote a trapezoidal fuzzy number as $\text{TrFN}(a,b,c,d)$. In the special case where $b=c$, the core collapses into a single value and the resulting TrFN is known as a *triangular fuzzy number*, or TFN. To facilitate ease of computation when dealing fuzzy numbers specified in general form, trapezoidal approximations of GLRFNs can be obtained using the methods described by Grzegorzewski and Mrowka [15,16].

Given a set of linguistic variables established for a specific problem, the set of possible fuzzy values each can take must be sensitive to individual interpretation and dimensional precision. As noted by Wallsten and Budescu [17], the location, spread, and shape of membership functions vary over individuals and depend upon context and the intent of the communicated message. Moreover, intra-individual vagueness in meaning must also be considered due to individual differences in understanding and operational lexicons, which has led researchers to suggest that words cannot be legislated [17]. It is therefore important to formally elicit membership functions that are specific to each problem or variable, and if resources permit, for each user using established techniques for expert opinion elicitation [18]. One effective technique is the Multistimuli Membership Function Technique (MMFT) described by Budescu et al. [19].




Figure 1. Fuzzy numbers for selected probability words

3.2. Constructing a Fuzzy System

A fuzzy system is a collection of “if-then” rules that link a string input of linguistic variables (i.e., the antecedent) to an output value (i.e., the consequent) [20]. More specifically, a fuzzy system F approximates the true function $Y = f(X_1, X_2, \dots, X_M)$ via a set of N fuzzy inference rules defined on the input-output state space $X_1 \times X_2 \times \dots \times X_M \times Y$ of the form (shown for two input variables):

$$\text{if } (X_1 \text{ is } A_k) \text{ and } (X_2 \text{ is } B_k) \text{ then } (Y \text{ is } C_k) \quad (6)$$

where $A_k, B_k,$ and C_k are linguistic or fuzzy values assigned to $X_1, X_2,$ and $Y,$ respectively for the k^{th} rule ($1 < k \leq N$). The approximation of a true function f by a fuzzy system F is achieved via a set of overlapping fuzzy rule patches such as is shown in Figure 2 that cover part of the graph of an unknown or unascertained function. If each X_j ($1 < j \leq M$) can assume one of n_j values, the total number of rules N is:

$$N = \prod_{j=1}^M n_j \quad (7)$$

Accordingly, a higher resolution approximation with greater coverage of an unknown or unascertained function follows from a larger set of inference rules. For illustration, a fuzzy system consisting of 6 input variables each with three possible states is characterized by $3^6 = 729$ rules. If each variable could take on one of four possible states, the fuzzy system would then be characterized by $4^6 = 4,096$ rules. Thus, as the number of possible states for each linguistic variable increases, so too does the required number of inference rules; this is known as the “curse of dimensionality” [20]. The challenge is to balance the need for analytical resolution with simplicity, which thus requires a trade-off between the number of input states permitted for each linguistic variable in a fuzzy inference rule and the decision makers’ tolerance for precision. Methods for negotiating this tradeoff, though important, are outside the scope of this paper.




Figure 2. Approximation of a function $Y = f(x)$ with a set of fuzzy rule patches (trapezoidal and triangular membership functions shown on the X and Y axes)

Constructing a fuzzy system requires systematically evaluating the consequent or output value for each of N antecedents to generate a complete set of fuzzy inference rules. This can be done manually by a panel of experts using opinion elicitation techniques [18], or if empirical data is available, can be done using numerical techniques [21]. While the most simple and primitive approach to establishing a complete set of fuzzy inference rules is to use brute-force to

evaluate the consequent for each individual antecedent in turn, this method can be time consuming if the number of rules is large or cognitively prohibitive if the number of input variables exceeds the mental abilities of a group of experts to process them into an opinion [22]. Fortunately, the fuzzy control literature is replete with suggestions on efficiently obtaining and updating a complete set of fuzzy inference rules, such as the recent article by Kuo et al. [23] using genetic algorithm-based fuzzy neural networks or other suggestions in [20]. The choice of which elicitation technique to use is at the discretion of the analyst.

3.3. Fuzzification and Rule Matching

Given a fuzzy system F consisting of a set of N inference rules that map input variables X_j ($1 < j \leq M$) to an output Y , a series of crisp input values x_j can be processed through the rule base to obtain an approximation of y with associated uncertainty. In particular, given a set of crisp input values x_j , the membership values $m(x_j)$ for all linguistic values or fuzzy numbers that can be assigned to X_j are determined. This is known as fuzzification, or the process of encoding a crisp number for use with a fuzzy rule base. For example, given a probability value $p=0.5$, the corresponding degree of membership for each fuzzy number in Figure 1 is $m_{\text{Tossup}}(0.5) = 1$, $m_{\text{Likely}}(0.5) = 0.5$, $m_{\text{Very Likely}}(0.5) \cong 0.1$, and $m(0.5) = 0$ for all others. Thus, if the linguistic variable X represents a probability, the fact that the $p = 0.5$ means that X is “Tossup” to the degree 1.0, X is “Likely” to the degree 0.5, and so forth.

Once the membership values for the fuzzy numbers evaluated at x_j are obtained, the next step is to determine which of the N rules are relevant (i.e., which rules are “on” or are “firing”) with respect to the given the input state. That is, the rule matching step assigns a degree of membership to each inference rule according to the degree of membership of its premises associated with x_j . Considering the example in the preceding paragraph, the rule “If X is ‘Likely’ then Y is ‘Bad’” would have a membership of 0.5 since $m_{\text{Likely}}(0.5) \cong 0.5$, whereas the rule “If X is ‘Unlikely’ then Y is ‘Good’” would have membership of zero in the set of relevant rules since $m_{\text{Unlikely}}(0.5) \cong 0.0$. For a set of rules with M premises, the membership, $m_{\text{Rule},k}$, for the k^{th} rule is determined according to the inputs x_j as follows:

$$m_{\text{Rule},k}(x_1, x_2, \dots, x_M) = \prod_{j=1}^M m_{X_j}^k(x_j) \quad (8)$$

where $m_{X_j}^k$ is the membership function of the fuzzy number assigned to X_j according to rule k .

Alternative conjunctive operators (e.g., minimum, maximum) can be used in lieu of the product in Eq. 8 [20,24], the choice being at the discretion of the analyst.

Note that while the discussion in this section focused on crisp inputs, the information and uncertainty associated with the input values x_j can be represented in any suitable form for expressing quantitative information, such as a probability distribution, interval, possibility distribution, grey number, etc. The form of the output from a fuzzy system captures the types of uncertainty present in the inputs. Moreover, the inputs can reflect the opinion of a single expert or the consensus opinion of a group of experts. To accommodate divergent opinions of multiple conflicting experts, one must supplement this model with appropriate techniques to aggregate the analysis results, such as through the use of weighting factors attached to each opinion [18].

3.4. Fuzzy Inference

In this final step, a crisp output y corresponding to the input state defined by x_j is obtained by combining all rules for which $m_{\text{Rule}} > 0$ to obtain a fuzzy representation of Y , then converting the aggregate fuzzy output into a crisp value. More specifically, the aggregate fuzzy representation for Y , $b_Y(x)$, is obtained from the linear combination of fuzzy numbers representing the output for each rule, each weighted according to the membership of the corresponding rule k ($1 < k \leq N$) as determined from Eq. 8:

$$b_Y(x_1, x_2, \dots, x_M, y) = \sum_k^N m_{\text{Rule},k}(x_1, x_2, \dots, x_M) m_Y^k(y) \quad (9)$$

where $m_Y^k(y)$ is the membership function of the fuzzy number assigned to the output linguistic variable Y according to rule k . The crisp output for $y = F(x)$ (letting x imply x_1, x_2, \dots, x_M for brevity) is obtained through *defuzzification* of Y via the *center of gravity method* as follows [20]:

$$y = F(x) = \frac{\int_{-\infty}^{+\infty} u b_Y(x, u) du}{\int_{-\infty}^{+\infty} b_Y(x, y) du} \quad (10)$$

The fuzzy system obtained via the aggregation and defuzzification operations in Eqs. 9 and 10 is known as an *additive fuzzy system* [20]. Alternative methods for inference and defuzzification can be found in the literature e.g., [24].

The crisp value for y obtained from Eq. 10 can be interpreted as the expected value of the output y given the input state defined by x_j . Moreover, it can be shown [20] that by defining $p_k(x)$ as:

$$p_k(x) = \frac{m_{\text{Rule},k}(x) V_Y^k}{\sum_{i=1}^N m_{\text{Rule},i}(x) V_Y^i} \quad (11)$$

where $V_Y^k = \int_{-\infty}^{+\infty} m_Y^k(y) dy$ is the total area of under the curve associated with the membership function $m_Y^k(y)$, Eq. 10 can be rewritten as:

$$F(x) = \sum_{k=1}^M p_k(x) c_Y^k \quad (12)$$

where $c_Y^k = \int_{-\infty}^{+\infty} y m_Y^k(y) dy / V_Y^k$ is the centroid of the membership function $m_Y^k(y)$ along the y -axis.

The standard deviation of the output y given the input state defined by x_j , $\sigma_{Y|X}$, can now be expressed as:

$$\sigma_{Y|X}(x) = \sqrt{\sum_{k=1}^M p_k(x) \sigma_{Y,k}^2 + \sum_{k=1}^M p_k(x) (c_Y^k - F(x))^2} \quad (13)$$

where:

$$\sigma_{Y,k}^2 = \frac{1}{V_Y^k} \int_{-\infty}^{\infty} (y - c_Y^k)^2 m_Y^k(y) dy \quad (14)$$

The standard deviation in Eq. 13 captures the epistemic uncertainty induced by the imprecision of the fuzzy numbers used for defining the functional relationship between input and output variables. An alternative procedure for performing fuzzy inference using random sets and probability boxes has been described by McGill [26].

4. Methodology

In general, the effectiveness of a security system typically depends on the capabilities of the defender to detect an adversary, delay an adversary long enough to engage once detected, and defeat the adversary when engaged [1,2]. Based on a discussion with a team of security experts, Morgenson, et al. [25] identified a set of six defensive criteria that characterize the effectiveness of a target's (e.g., facility or asset) security capabilities as described in Table 1.

Table 1. Defensive criteria for probability of adversary success assessment [25]

Variable	Defensive Criterion	Definition
X ₁	Access control	Consists of the part of the security system focused on controlling access to the facility
X ₂	Personnel barriers	Provides a means of denying or delaying unauthorized personnel access into and within the facility
X ₃	Vehicle barriers	Provides a means of denying or delaying unauthorized vehicular access into and within the facility
X ₄	Surveillance systems	Consists of all measures to detect the presence of a threat, including intruders and individuals with malicious intent
X ₅	Guard force	Provides a means to engage and neutralize an identified threat; also serves as an additional barrier for personnel and vehicles and provides an additional means for detection
X ₆	Reaction force with heavy weapons	Provides a means to effectively engage and neutralize an identified threat

Applying the ideas presented in section 3, a fuzzy system can be constructed (Figure 3) that approximates the functional relationship between the subjective assessment of each defensive criterion to an output probability of adversary success via an exhaustive set of linguistic statements of the form:

if X_1 and X_2 and X_3 and X_4 and X_5 and X_6 , then $\Pr(S|A)$ (15)

where X_i are the premises from Table 1 specified as linguistic variables with membership functions constructed over a domain suitable for its measured (e.g., $[0,10]$, $[0,100]$, or $[0,1]$ on a constructed scale, $[0, \infty)$ on a scale for a measurable quantity, etc.), and $\Pr(S|A)$ is the consequent which is a linguistic variable specified over the domain $[0,1]$ as required by the axioms of probability theory [26]. For the purpose of further discussion and illustration, the premises X_i may take on linguistic values “Low,” “Medium,” or “High” defined on a constructed scale for effectiveness with membership functions shown in Figure 4, and the consequent $\Pr(S|A)$ may take on linguistic values such as “Likely,” “Certain,” or “Even Chance” with membership functions shown in Figure 5. In practice, the phrases and membership functions used in constructing a set of fuzzy inference rules would be elicited from a panel of experts as later described in section 6, and may be broadly applicable to all initiating events and all premises, or have meanings that are pegged to specific contexts or defined for specific defensive criteria. Note that the phrases for effectiveness in Figure 4 and probability in Figure 5 can, in principle, be used in the context of any attack type so long as the membership functions are calibrated to the user or group of users. However, it is permissible to use a specific set of phrases depending on the attack type; for example, some threats might warrant a higher resolution assessment in some criteria or require different metrics. It is also possible that each defensive criterion may require its own set of linguistic phrases for effectiveness, such as would be the case if one criterion was based on a constructed scale, another on a crisp scale such as time, etc.

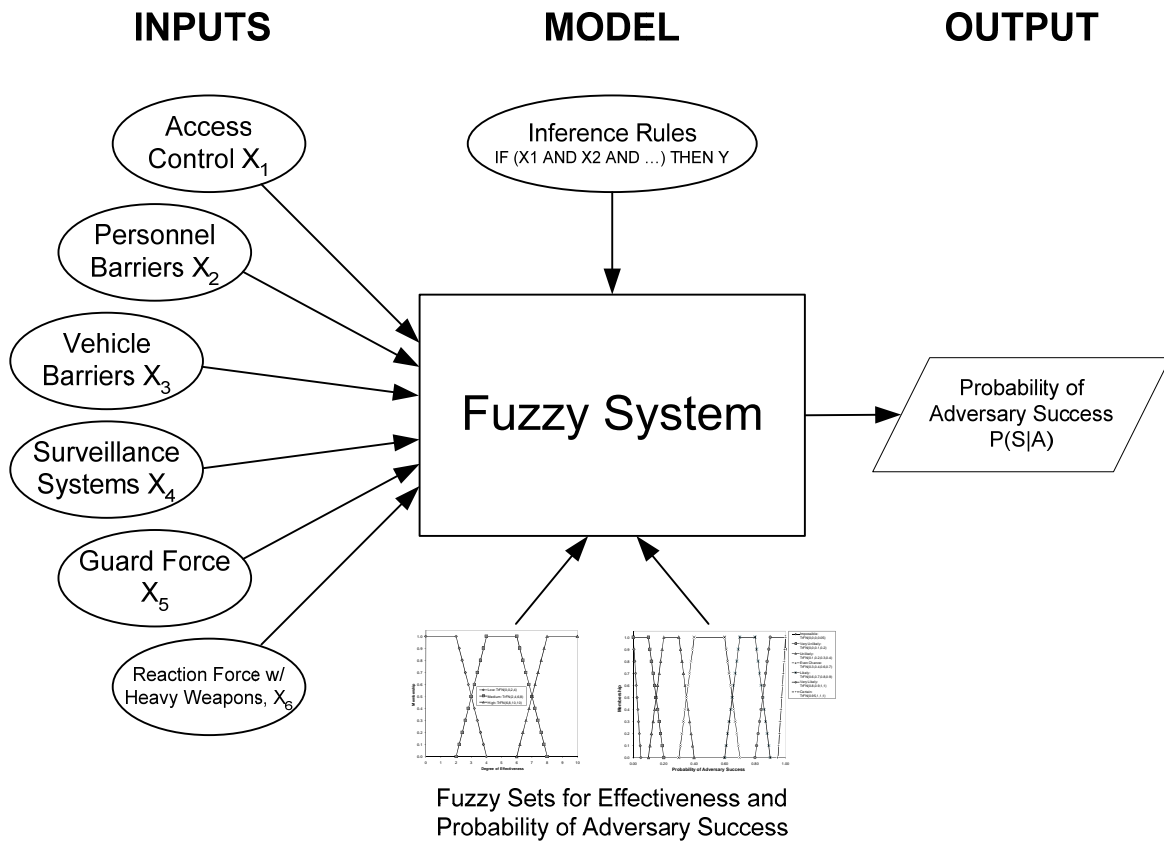


Figure 3. Schematic of the fuzzy system architecture for security system performance assessment




Figure 4. Notional membership functions for the fuzzy numbers representing degree of effectiveness on a constructed scale




Figure 5. Notional membership functions for some fuzzy numbers representing probability of adversary success

Given a set of premises X_i with prescribed linguistic values, the consequent $\Pr(S/A)$ is specified for each by a panel of security experts. The set of all possible combinations of linguistic values assigned to the premises represents an exhaustive set of fuzzy rules, which would contain 729 rules according to Eq. 7 using the three fuzzy numbers in Figure 4. For example, a panel of experts might decide that the following set of premises on the effectiveness of security measures corresponds to a probability of success equal to “Very Likely” (selected from the fuzzy numbers in Figure 5):

$$\begin{aligned}
 &\text{If } (X_1 \text{ is "Medium"}) \text{ and} \\
 &\quad (X_2 \text{ is "Low"}) \text{ and} \\
 &\quad (X_3 \text{ is "Low"}) \text{ and} \\
 &\quad (X_4 \text{ is "High"}) \text{ and} \\
 &\quad (X_5 \text{ is "Low"}) \text{ and} \\
 &\quad (X_6 \text{ is "Low"}) \\
 &\text{Then } (\Pr(S/A) \text{ is "Very Likely"})
 \end{aligned} \tag{16}$$

When combined, the complete set of fuzzy inference rules of this form approximate the true functional relationship between inputs and outputs, and thus provide an alternate means for specifying the function $\Pr(S/A) = f(X_1, X_2, \dots)$ without having to define an explicit mathematical

formula. Figure 6 illustrates an exhaustive set of fuzzy inference rules for a notional initiating event of the form in Eq. 15 using the fuzzy numbers shown in Figure 4 for the premises, where the rule number, Z , shown in each cell is obtained as:

$$Z = \sum_{i=1}^6 3^{i-1} \hat{X}_i \quad (16)$$

where \hat{X}_i takes the value 0, 1, or 2 according to whether the premises X_i is “Low,” “Medium,” or “High,” respectively, and the cell background shading and pattern maps to a probability phrase shown in the legend. For example, the rule described in the example of Eq. 16 would have $Z = (3^0)(1)+(3^1)(0)+(3^2)(0)+(3^3)(2)+(3^4)(0)+(3^5)(0) = 55$ and would be shaded with vertical dark and light grey stripes. Note that in general, a unique set of inference rules must be constructed for each initiating event considered (e.g., explosive attack with truck bomb or an explosive attack via an unmanned aerial vehicle). The reason for this is that the defensive criteria interact in different ways depending on the nature of the threat, and thereby affects the probability of adversary success.

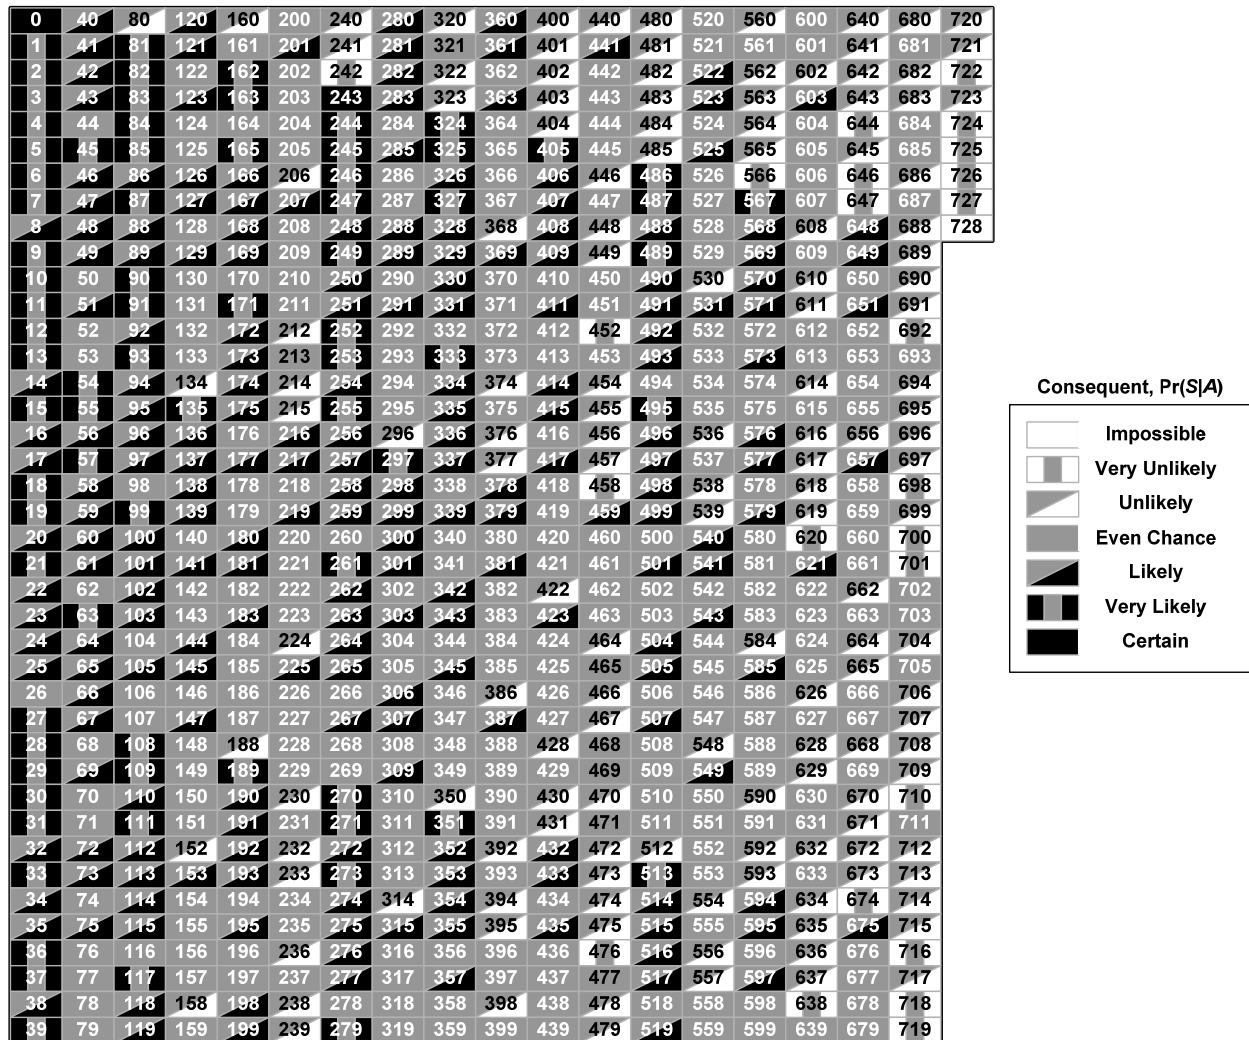


Figure 6. Exhaustive set of fuzzy inference rules

Once the fuzzy inference rules are defined, a user such as a security expert can subjectively assign a value to each premise or criterion on a scale of 0 to 10 (or on an alternate scale if preferred) for a given facility or asset and attack type. Note that this elicitation step is different from that used to construct the underlying fuzzy logic model. According to the proposed model, the corresponding membership values for each relevant fuzzy set associated with the premises would be obtained, processed according to the set of fuzzy inference rules, and then the result would be translated back into a crisp value for probability of adversary success (mean) with an associated expression of epistemic uncertainty (standard deviation) via defuzzification as described in sections 3.3 and 3.4. Given the constraints on the probability of adversary success defined by the output mean, standard deviation, and the axioms of probability, maximum entropy arguments insist that the epistemic uncertainty on the true probability of adversary success be represented by a beta distribution over a domain $[a,d]$ (see section 3.1) spanning the support of the aggregate fuzzy set [12]. This distribution for $\Pr(S|A)$ can now be used in conjunction with the probability of attack and loss given success in Eq. 2 to estimate the risk to an asset or facility.

Again, note that the goal of the proposed model is to provide a system based on approximate reasoning that produces an estimate for the probability of adversary success based on the subjective evaluation of several or more defensive criteria. As described in section 3.4, this model produces a point estimate of the probability of adversary success and a standard deviation about this point estimate that characterizes the epistemic uncertainty. In theory, one can determine an error rate or error by comparing the results from the proposed model (which are probabilistic in two senses) with another value of the probability estimated by other means. The discussion in section 6.3 offers some suggestions on how to do this. Furthermore, as with any model, it should be noted that the results are only as good as the inputs used to generate them. If two or more experts use the same model with the same inference rules to look at the same facility and initiating event, there should be no expectation that the model will produce the same result if each expert assesses things differently (though we would hope divergent opinions would prompt discussion and debate).

5. Illustrative Example

To demonstrate the proposed approach to security system performance assessment in an operational environment, a prototype of the methodology described in the preceding sections was implemented in a Microsoft Excel 2007 environment with user interface shown in Figure 7. A series of Excel worksheets were constructed for collecting expert opinions on the security posture of a facility according to the six defensive criteria (e.g., Figure 7), storing the fuzzy numbers for degree of effectiveness for each defensive criterion and probability of adversary success (e.g., Figures 3 and 4), specifying the set of fuzzy inference rules (e.g., Figure 5), and performing fuzzy inference (fuzzification, rule matching, and defuzzification) according to procedures outlined section 3.4. The spreadsheet tool can be made available for academic purposes only to interested readers upon written request to the authors. For the purposes of demonstration, this example leverages the fuzzy system defined according to Eq. 15 and the inference rules in Figure 6 for a notional initiating event defined according to the fuzzy numbers shown in Figures 3 and 4. As shown in the Figure 7 for this notional initiating event, a group of security experts reached a consensus assessment for each defensive criterion as follows: $X_1 = 2.9$, $X_2 = 4.6$, $X_3 = 6.0$, $X_4 = 3.5$, $X_5 = 7.9$, and $X_6 = 4.3$. Note that this group could just as easily have expressed their assessment as intervals, probability distributions, etc. Table 2 provides the

corresponding membership values for each fuzzy number associated with the defensive criteria, and Table 3 summarizes the corresponding list of active fuzzy rules with associated consequents and weights. According to Table 3, a total weight of 0.164 is assigned to the output fuzzy set “Even Chance” and a total weight of 0.836 is assigned to the output set “Unlikely” (weights to all other output fuzzy sets are zero). The aggregate output set based on a weighted sum of the two active output fuzzy sets was determined from Eq. 8 and is shown in Figure 8. Using the center of gravity method described in Eqs. 12 and 13, the defuzzified result from this assessment yields a probability of adversary success with a mean of 0.53 and standard deviation of 0.12 (0.53 is the center of gravity of the output fuzzy set along the horizontal axis). Using maximum entropy arguments [12], given constraints on mean, standard deviation, and the domain of the distribution defined by the support of the aggregate output set, the epistemic uncertainty on the output probability of adversary success can be represented by a beta distribution over the domain (0.3,0.9) shown as a solid line in Figure 8.

Table 2. Memberships in each fuzzy set corresponding to the assessed values for the defensive criteria

Defensive Criterion	Value	Membership Values in Fuzzy Sets		
		“Low”	“Medium”	“High”
X_1	2.9	0.55	0.45	0.00
X_2	4.6	0.00	1.00	0.00
X_3	6.0	0.00	1.00	0.00
X_4	3.5	0.25	0.75	0.00
X_5	7.9	0.00	0.05	0.95
X_6	4.3	0.00	1.00	0.00

Table 3. Active rules, consequents, and weights

Active Rule (Z)	Consequent Pr(S A)	Weight
336	Even Chance	0.007
337	Even Chance	0.131
363	Even Chance	0.021
364	Unlikely	0.392
417	Even Chance	0.006
418	Unlikely	0.107
444	Unlikely	0.017
445	Unlikely	0.321

Multicriteria Security System Performance Assessment Tool

Defensive Criterion	Effectiveness Assessment		Reset Worksheet
	(Values correspond to the degree of effectiveness for each criterion, where 0 = completely ineffective and 10 = perfect)		
X ₁ : Access Control			2.9
X ₂ : Personnel Barriers			4.6
X ₃ : Vehicle Barriers			6.0
X ₄ : Surveillance Systems			3.5
X ₅ : Guard Force			7.9
X ₆ : Reaction Force with Heavy Weapons			4.3

Probability of Adversary Success		
Pr(S A) =	Mean	Standard Deviation
	0.53	0.12

Figure 7. User interface for the Microsoft Excel security system performance assessment tool; the sliders are used to assign values to each of the six defensive criteria shown in the right-most column




Figure 8. Aggregate output set for defuzzification and associated maximum entropy distribution

6. Implementation

The successful implementation of the ideas presented so far in an operational setting should follow a three phase process as suggested in Figure 9. Each phase is described as follows.

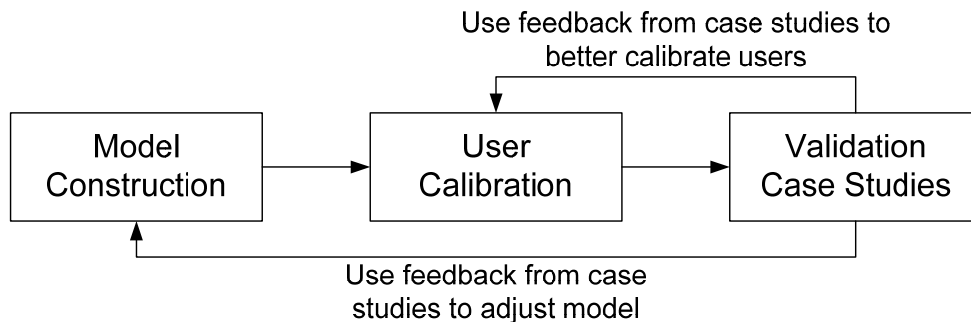


Figure 9. Suggested process for implementing the proposed security evaluation model

6.1. Model Construction

This phase constructs a logical model that approximates the functional relationship among several defensive assessment criteria to estimate the probability of adversary success for a

variety of initiating events and attack profiles. More specifically, a set of defensive criteria variables such as the ones described in [25] are established with associated linguistic values for degree of effectiveness (e.g., “Low,” “Medium,” and “High”). An exhaustive set of linguistic rules of the form in Eq. 15 is constructed for the fuzzy inference rules by establishing all possible combinations among fuzzy values for each defensive criterion, resulting in a total number of rules as defined by Eq. 7. For each antecedent, a team of qualified security experts assigns a unique linguistic value to the probability of adversary success for a variety of initiating events and attack profiles using a suitable elicitation process for constructing fuzzy rules. The linguistic values for probability of adversary success can be chosen from a prescribed list of words or elicited directly from the experts. The antecedents combined with a corresponding probability assignment define the set of fuzzy inference rules for a specific initiating event, and yields an approximate model of the relationship between input criteria and probability of adversary success for that event.

6.2. Expert/Team Calibration

This phase calibrates the quantified judgment of a security assessment team with the linguistic scales for the input variables identified in the previous step using proven elicitation techniques (e.g., [12,19]). For each linguistic value characterizing degree of effectiveness and probability of adversary success identified in the previous phase, a membership function is obtained that maps to it a range of numeric values on a suitable numeric scale, such as 0 to 10 for a subjective scale, [0,1] for probabilities, or in native units if a criterion represents a physically measurable quantity (e.g., time to detect as measured in seconds). A goal should be to develop sufficient guidance on the meaning of each criterion so as to minimize variation in interpretation and understanding. Moreover, if it is desired to aggregate the assessment of multiple experts, weighting factors can be assigned to each expert according to his expertise in relation to assessing each criterion. The computational details for this feature, however, are left to the reader to explore.

6.3. Model Validation

This phase attempts to validate the constructed security evaluation model with respect to benchmark assessments of probability of adversary success for a series of case studies. The results from this model would be compared by with results obtained through other means, such as through probabilistic modeling using the approach described in [5], to determine whether the outputs make sense respect to the inputs. Due to the complexity of the security effectiveness assessment problem, it is anticipated that suitability of the model would ultimately be determined by a panel of human experts; however, it is left to the reader to decide the appropriate model validation approach. The results from this phase can also be used to adjust or correct the fuzzy inference model or better calibrate the users.

7. Conclusions

The performance of a security system under load is generally difficult and expensive to measure, often requiring significant time and analytical resources to make defensible quantitative statements about security. The primary challenges to quantification of security system performance, usually expressed in terms of probability of adversary success for a suite of initiating events or attack profiles, are the complexity of the security system and all its hardware,

software, and human elements combined with the highly uncertain nature of an adaptive and intelligence adversarial loading environment [1,2]. This complexity often clashes with the need to produce rapid assessments of overall security risk. For example, the allocation of federal government resources among urban areas for national risk reduction over the course of a fiscal year requires quick, yet defensible estimates of urban area risk, which in turn requires information on the risks associated with those assets and systems situated in the urban area. If one considers that there are tens of thousands of assets, each taking several days (if not more) to assess using conventional methods, the magnitude of the resources required to obtain risk information within time frame to support fiscal decisions (i.e., 1 year) make conventional methods prohibitive. Thus, simple techniques are needed to assess security system performance that are both acceptable to decision maker and defensible in the face of public scrutiny.

This paper introduced a simple model for assessing the probability of adversary success for a given initiating event based on the subjective evaluation of several or more security effectiveness or defensive criteria. Based on proven techniques of fuzzy logic, the proposed approach is “model-free” in the sense that a functional relationship between input and output values is not mathematically explicit, but rather emerges from a set of fuzzy inference rules of the form “If X then Y ” that consider all possible combinations of linguistic values that can be assigned to the input linguistic variables X_i . Once developed, the proposed model is quick to implement provided the security experts are sufficiently trained to interpret facility-specific data and information in relation to properly defined evaluation criteria. Moreover, since the output from the proposed model represents a probability (with associated uncertainty on the value), the results are compatible with the data needs for quantitative security risk assessment, such as was presented in Eq. 2. It is expected that implementation of this model in the homeland security community will provide an effective means of quickly obtaining useful information to support risk-based resource allocation decisions.

8. References

1. Ayyub, B. M., McGill, W. L., and Kaminskiy, M. P. (2007). “Critical Asset and Portfolio Risk Analysis: An All-Hazards Framework.” *Risk Analysis*, 27(4): 789-801.
2. McGill, W. L., Ayyub, B. M., and Kaminskiy, M. P. (2007). “Risk Analysis for Critical Asset Protection.” *Risk Analysis*, 27(5): 1265-1281.
3. Ayyub, B. M. (2003). *Risk Analysis in Engineering and Economics*, CRC Press
4. Kaplan, S., and Garrick, B. J. (1981). “On the Quantitative Definition of Risk.” *Risk Analysis*, 1(1): 11-27.
5. Dessent, G. H. (1987). “Prison Perimeter Cost Effectiveness.” *Journal of the Operational Research Society*, 10: 975-980.
6. Hicks, M. J., Snell, M. S., Sandoval, J. S., and Potter, C. S. (1999). “Physical Protection Systems – Cost and Performance Analysis: A Case Study.” *IEEE AES Systems Magazine*, April 1999.
7. Matalucci, R. V. (2002). “Risk Assessment Methodology for Dams (RAM-D).” *Proceedings of the 6th International Conference on Probabilistic Safety Assessment and Management (PSAM6)*, 23-28 June 2002, San Juan, Puerto Rico, USA, Vol. 1, 169-176.

8. Kaplan, S., Visnepolshi, S., Zlotin, B., and Zusman, A. (1999). *New Tools for Failure & Risk Analysis: Anticipatory Failure Determination (AFD) & the Theory of Scenario Structuring*. Ideation International.
9. Manunta, G. (1999b). "What is Security?" *Security Journal*, 12(3): 57-66.
10. NIICIE (2007). "CARVER2®: Critical Infrastructure Assessment Tool." National Infrastructure Institute Center for Infrastructure Expertise. Available at: <http://www.ni2ciel.org/CARVER2.asp>.
11. Zadeh, L. (1975). "The Concept of a Linguistic Variable and its Application to Approximate Reasoning, Part I." *Information Sciences*, 8(3): 199-249.
12. Ayyub, B. M., and Klir, G. J. (2007). *Uncertainty Modeling and Analysis in Engineering and the Sciences*, Chapman & Hall/CRC Press.
13. Lichtenstein, S., and Newman, J. R. (1967). "Empirical Scaling of Common Verbal Phrases Associated with Numerical Probabilities." *Psychonomic Science*, 9(10): 563-564.
14. Dubois, D., and Prade, H. (1980). *Fuzzy Sets and Systems: Theory and Applications*. Elsevier Science & Technology Books.
15. Grzegorzewski, P., and Mrowka, E. (2005). "Trapezoidal Approximations of Fuzzy Numbers." *Fuzzy Sets and Systems*, 153(1): 115-135
16. Grzegorzewski, P., and Mrowka, E. (2007). "Trapezoidal Approximations of Fuzzy Numbers - Revisited." *Fuzzy Sets and Systems*, 158(7): 757-768.
17. Wallsten, T. S., and Budescu, D. V. (1994). "A Review of Human Linguistic Probability Processing: General Principles and Empirical Evidence." *The Knowledge Engineering Review*, 10(1): 43-62.
18. Ayyub, B. M. (2001). *Elicitation of Expert Opinions for Uncertainty and Risks*. CRC Press.
19. Budescu, D. V., Karelitz, T. M., and Wallsten, T. S. (2003). "Predicting the Directionality of Probability Words from their Membership Functions." *Journal of Behavioral Decision Making*, 16(3): 159-180.
20. Kosko, B. (1997). *Fuzzy Engineering*. Prentice Hall: Upper Saddle River, NJ.
21. Chiu, S. L. (1999). "Extracting Fuzzy Rules from Data for Function Approximation and Pattern Classification." In Dubois, D., Prade, H., and Yager, R. eds. (1997). *Fuzzy Information Engineering: A Guided Tour of Applications*. Wiley & Sons.
22. Miller, G. A. (1956). "The Magical Number Seven, Plus or Minus Two: Some Limits On Our Capacity for Processing Information." *The Psychological Review*, 63(2): 81-97.
23. Kuo, R. J., Hong, S. M., Sheu, J.-B., Lin, Y., and Huang, Y. C. (2007). "Continuous Genetic Algorithm-Based Fuzzy Neural Network for Learning Fuzzy IF-THEN Rules." *Neurocomputing*, Article In Press.
24. Passino, K. M., and Yurkovich, S. (1998). *Fuzzy Control*. Addison-Wesley: Menlo Park, CA.
25. Morgeson, J. D., Utgoff, V. A., Fainberg, M. A., and Keleher, M. (2006). "National Comparative Risk Assessment Pilot Project. Volume I: Main Text with Appendixes A and B." IDA Document D-3309, Draft Final, Institute for Defense Analyses: Alexandria, VA.
26. Ayyub, B. M. and McCuen, R. H. (2002). *Probability, Statistics, and Reliability for Engineers and Scientists*. Chapman & Hall/CRC Press.
27. McGill, W. L. (2008). *Critical Asset and Portfolio Risk Analysis for Homeland Security*, PhD Dissertation, University of Maryland [in preparation].

Author Biographies

William L. McGill is an assistant professor of information sciences and technology at the Pennsylvania State University (University Park) where he is a member of the security risk analysis faculty. He is author or coauthor of numerous reports and publications in journals and conference proceedings centered on risk analysis for safety, security, and intelligence. His research interests are in the development and application of innovative risk analysis methodologies in support of security and defense decision making.

Bilal M. Ayyub is a professor of civil and environmental engineering at the University of Maryland (College Park) and the director of the Center for Technology and Systems Management. Dr. Ayyub is a fellow of ASCE, ASME and SNAME, and has served the engineering community in various capacities through societies that include ASNE, ASCE, ASME, SNAME, IEEE-CS, and NAFIPS. He is the author and co-author of more than 500 publications in journals and conference proceedings, and reports. His publications include several textbooks. Dr. Ayyub is a multiple recipient of the ASNE "Jimmie" Hamilton Award for the best papers in the Naval Engineers Journal in 1985, 1992, 2000 and 2003. Also, he received the ASCE "Outstanding Research Oriented Paper" in the Journal of Water Resources Planning and Management for 1987, the ASCE Edmund Friedman Award in 1989, the ASCE Walter Huber Research Prize in 1997, and the K. S. Fu Award of NAFIPS in 1995.