LA-UR-03-3467

*Approved for public release;
distribution is unlimited.*

Title:

Assessing the Risk of Nuclear Terrorism Using Logic Evolved Decision Analysis

Author(s):

Steve Eisenhawer
Terry Bott
D.V. Rao

Submitted to:

American Nuclear Society Annual Meeting
June 1-4, 2003
San Diego, CA

Assessing the Risk of Nuclear Terrorism Using Logic Evolved Decision Analysis

S. W. Eisenhawer
Los Alamos National Laboratory, MS K557, Los Alamos, New Mexico 87545 seisenhawer@lanl.gov
T. F. Bott
Los Alamos National Laboratory, MS K557, Los Alamos, New Mexico 87545 tbott@lanl.gov
D. V. Rao
Los Alamos National Laboratory, MS F606, Los Alamos, New Mexico 87545 dvrao@lanl.gov

Abstract–*A necessary first step in the allocation of resources to increase homeland security is an evaluation of the risk of terrorist attack. In this paper we demonstrate the use of Logic Evolved Decision analysis (LED) to the estimation of the risk of nuclear terrorism. LED uses linked logic models to represent the elements of the decision process. These logic models are built using the software tool LED TREE. The first model, the possibility tree is used to deduce a set of attack scenarios. The second model, the inference tree represents the evaluation model used to infer risk for the individual attack scenarios. The model proposed here is based on a game theoretic perspective where the set of attackers and the defender play an extensive game with imperfect information. We perform the risk evaluation using approximate reasoning (AR). AR uses a series of forward-chained rule bases to emulate expert judgment. It is particularly well suited to decision problems where much of the data is qualitative and many of the relevant factors and their importance are perceptual in nature. Results for an illustrative problem with a small set of attack scenarios are presented.*

I. INTRODUCTION

Following the events of September 11, 2001, there has been a quantum change in the level of attention directed toward defending the homeland against terrorist attack. The scope of current U. S. counter-terrorism efforts is reflective of the spectrum of possible threats that require consideration. It is clearly impossible to protect against all possible threats and an approach to prioritization is needed. A natural metric to use in the allocation of resources to this problem is risk. However, having chosen risk as a suitable metric, a number of important issues remain to be resolved

- What is a comprehensive set of threat scenarios for which risk is to be estimated?
- What is the likelihood that a particular scenario will be attempted?
- Given an attempt what is the likelihood that the attempt will be successful?
- How are the diverse consequences associated with a scenario to be suitably aggregated?
- How should a risk ranking of the scenarios be presented and what is the appropriate measure to express the confidence in the results?

These and other related questions are similar to those confronted in other risk analyses and one could easily reach the conclusion that an adaptation of standard PSA techniques is called for. We argue here that a different approach is better suited to the largely qualitative knowledge that exists in the counter terrorism field. The nature of this argument is outlined by considering a specific problem – nuclear terror.

The methodology presented here to assess terrorist risk is called Logic Evolved Decision (LED) analysis.[1, 2] The fundamental idea upon which LED is based is the use of linked logic models to represent the necessary functions of a decision analysis tool. Each logic model is a directed graph called a process tree. [3,4] The process trees are developed deductively using general-purpose tree construction software called LED TREE. [5] Two process trees are essential for decision analysis: a *possibility* tree that represents a comprehensive set of alternatives, in this case terrorist attack scenarios and an *inference* tree that defines how a metric is to be inferred. Here the metric is risk and this measure will be used to rank order the scenarios obtained from the possibility tree.

II. POSSIBILITY TREE FOR ATTACK SCENARIOS

Figure 1 shows the possibility tree for nuclear terror attacks constructed using LED TREE.[a] The intent in building this tree is to deduce a comprehensive set of

[a] We assume familiarity on the reader's part with standard logic gates and the hierarchical structure of a deductive tree such as a fault tree.

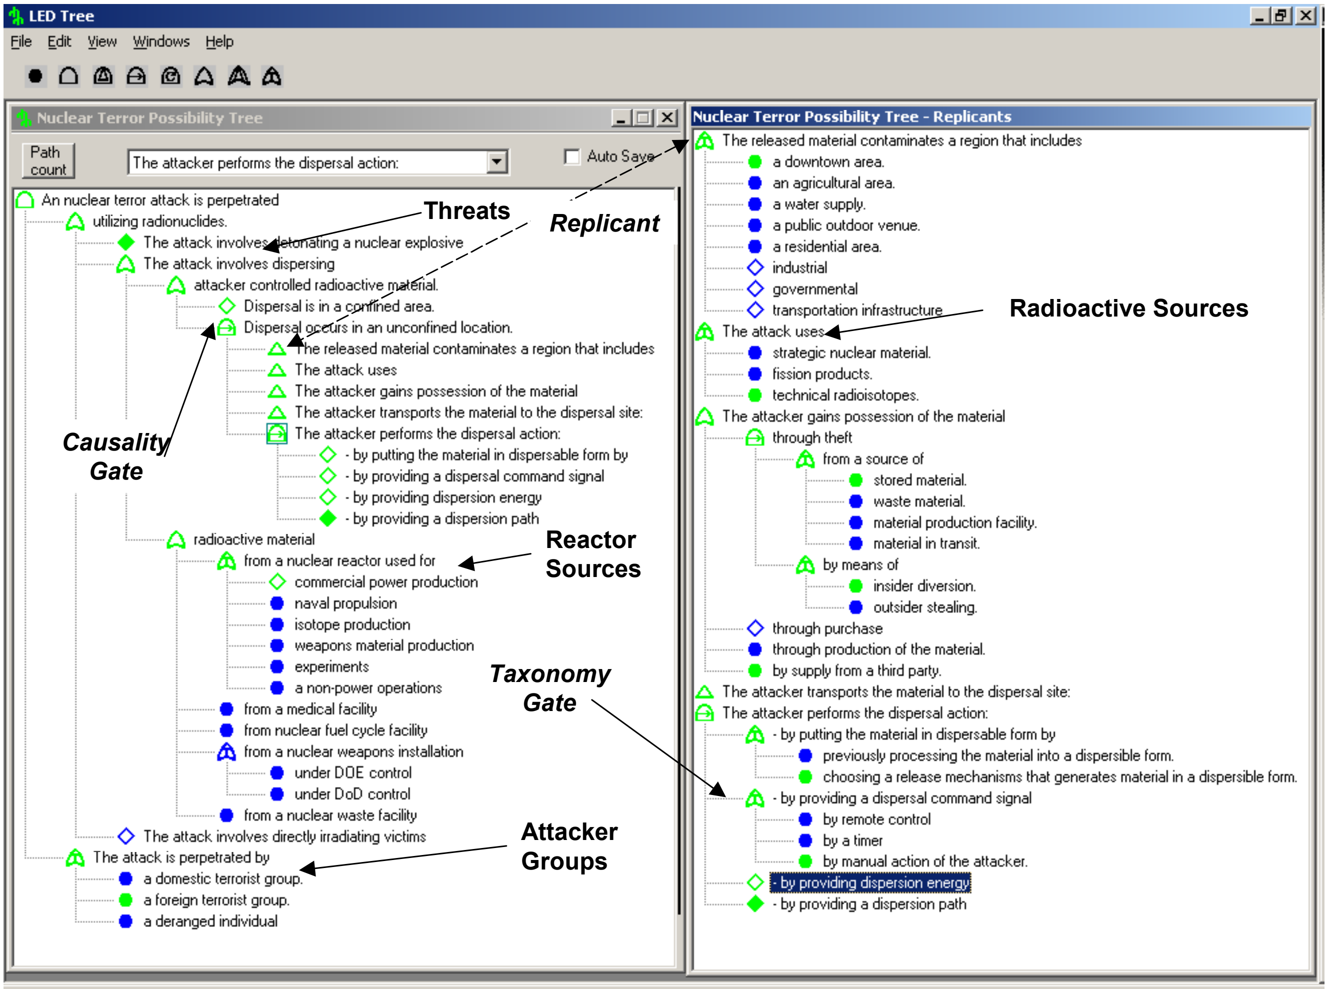

LED Tree
File Edit View Windows Help
Nuclear Terror Possibility Tree
Path count
The attacker performs the dispersal action:
Auto Save
Nuclear Terror Possibility Tree - Replicants

```
An nuclear terror attack is perpetrated
    utilizing radionuclides.
        The attack involves detonating a nuclear explosive
        The attack involves dispersing
            attacker controlled radioactive material.
                Dispersal is in a confined area.
                Dispersal occurs in an unconfined location.
                    The released material contaminates a region that includes
                    The attack uses
                    The attacker gains possession of the material
                    The attacker transports the material to the dispersal site:
                    The attacker performs the dispersal action:
                        - by putting the material in dispersable form by
                        - by providing a dispersal command signal
                        - by providing dispersion energy
                        - by providing a dispersion path
            radioactive material
                from a nuclear reactor used for
                    commercial power production
                    naval propulsion
                    isotope production
                    weapons material production
                    experiments
                    a non-power operations
                from a medical facility
                from nuclear fuel cycle facility
                from a nuclear weapons installation
                    under DOE control
                    under DoD control
                from a nuclear waste facility
        The attack involves directly irradiating victims
    The attack is perpetrated by
        a domestic terrorist group.
        a foreign terrorist group.
        a deranged individual
```

```
The released material contaminates a region that includes
    a downtown area.
    an agricultural area.
    a water supply.
    a public outdoor venue.
    a residential area.
    industrial
    governmental
    transportation infrastructure
The attack uses
    strategic nuclear material.
    fission products.
    technical radioisotopes.
The attacker gains possession of the material
    through theft
        from a source of
            stored material.
            waste material.
            material production facility.
            material in transit.
        by means of
            insider diversion.
            outsider stealing.
    through purchase
    through production of the material.
    by supply from a third party.
The attacker transports the material to the dispersal site:
The attacker performs the dispersal action:
    - by putting the material in dispersable form by
        previously processing the material into a dispersible form.
        choosing a release mechanisms that generates material in a dispersible form.
    - by providing a dispersal command signal
        by remote control
        by a timer
        by manual action of the attacker.
    - by providing dispersion energy
    - by providing a dispersion path
```

Threats

Replicant

Radioactive Sources

Causality Gate

Reactor Sources

Taxonomy Gate

Attacker Groups

Fig. 1. Possibility tree for attack scenarios.

scenarios involving nuclear weapons, radionuclide sources and aerosols from reactors. The top node in the tree shown to the left of the figure is "*A nuclear terror attack is perpetrated.*" This node is an AND gate, with inputs "*utilizing radionuclides*" and "*The attack is perpetrated by.*" The former is the start of the development of how the attack is carried out and the latter begins the description of who carries out the attack. The tree here is meant to be illustrative only. Representative attackers are suggested by the three terminal nodes "*a domestic terrorist group,*" "*a foreign terrorist group,*" and "*a disaffected individual.*" In an actual application each of these would be developed in more detail. The logic gate for "*The attack is perpetrated by*" is a TAXONOMY gate; LED TREE makes available a suite of special purpose gates so that the analyst can express quite specific logical relationships. Note also that the input nodes to this gate are in two different colors. The terminal node "*a foreign terrorist group*" is green (light tone), indicating that it is an active part of the logic equation. The other two terminals are blue (dark tone). These nodes have been designated as "excluded." Excluded status means that they are not currently part of the logic equation represented by the tree as shown; the only attacker group under consideration is "*a foreign terrorist group.*"

The node "*utilizing radionuclides*" is an OR gate. The first input to this gate "*The attack involves detonating a nuclear explosive*" appears as a solid green diamond. The diamond indicates that there is additional development that has been collapsed in this view. Collapsing a portion of the tree allows an analyst to concentrate on a particular section of the tree. The diamond here is solid, signifying that the collapsed sub tree has been "terminated." That is, the current version of the logic equation does not take into account the inputs to this node. The other inputs to "*utilizing radionuclides*" examine the possibilities associated with 1) the dispersal of radionuclides that the attacker controls, 2) aerosols from reactors and other components of the nuclear fuel cycle and 3) direct irradiation. We examine the first of these in more detail here. Logically the dispersal of attacker-controlled material could occur in a confined or unconfined location. The node "*Dispersal occurs in an unconfined location,*" is represented as a CAUSALITY gate indicating a process sequence. The first input to this gate "*The released material contaminates a region that includes*" appears as a triangle indicating a collapsed replicant. A replicant is a sub tree that can be used multiple times. All replicants are edited in the replicant window that appears to the right in Fig. 1. The remaining inputs to the gate - also replicants, describe the rest of the sequence resulting in the dispersal.

Our tree is a logical equation written in natural language form. One class of solutions is the set of paths. Each path is a unique attack scenario. For example:

> *A nuclear terror attack is perpetrated utilizing radionuclides. The attack involves dispersing attacker controlled radioactive material. Dispersal occurs in an unconfined location. The released material contaminates a region that includes a downtown area. The attack uses technical radioisotopes. The attacker gains possession of the material through theft from a source of stored material. by means of insider diversion. The attacker transports the material to the dispersal site: - using a road vehicle. - by placing the material in the transport vehicle. - carrying the material to the dispersal site. - gaining access to the dispersal site. The attacker performs the dispersal action: - by putting the material in dispersable form by choosing a release mechanisms that generates material in a dispersible form. - by providing a dispersal command signal by manual action of the attacker. - by providing dispersion energy through an explosion. - by providing a dispersion path. The attack is perpetrated by a foreign terrorist group.*

This scenario is unedited, that is it appears as shown from solving the logical equation. The relative ease in reading and understanding the scenarios demonstrates the value of using natural language in the possibility tree.

The number of paths and the degree of detail expressed in the individual paths is controlled by the use of the termination and exclusion commands. For the tree as shown in Fig. 1, there are six unique paths.[b] These are summarized in Table I. There are four attack scenarios involving radionuclide dispersal that differ according to target – an indoor sports arena, and an unspecified down-town area, and according to the method used to acquire the material – theft and supply from a third party. The third entry in the table corresponds to the path listed above. The fifth entry in the table is an attack on a nuclear reactor:

> *A nuclear terror attack is perpetrated utilizing radionuclides. The attack involves dispersing radioactive material from a nuclear reactor used for commercial power production: - by releasing radionculides from the reactor core by causing a core melt by causing a LOCA. The radioactive material is dispersed by natural dispersive phenomena. The released material contaminates a region that includes a downtown area. The attack is perpetrated by a foreign terrorist group.*

Finally the last entry is for an improvised nuclear device. The details for this attack appear in the solution when the termination of the node "*The attack involves detonating a nuclear explosive* " is removed:

> *A nuclear terror attack is perpetrated utilizing radionuclides. The attack involves detonating a nuclear explosive improvised by the attacker. The attack is perpetrated by a foreign terrorist group.*

We will consider this scenario in more detail in Section V.

III. INFERENCE TREE FOR RISK

Given a set of terrorist attack scenarios we would like to rank order them according to risk. Our inferential model incorporates a game theoretic perspective. The game to be played is asymmetric. A specific attacker will choose to attempt only a particular subset of attack scenarios associated with particular targets and employing specific attack modes. He will attempt to allocate his assets in order to inflict the maximum amount of terror. The defender on the other hand must try to protect all of the targets for which he bears responsibility against all attack scenarios. He will attempt to minimize his risk. This is an example of an extensive game that will be played with imperfect knowledge by all players.[6] We cast ourselves as the defender in this game and therefore wish to minimize attack risk. To achieve this we must take into account the counter-objective of the attackers. That is, in order to estimate risk, we must make an estimate of the optimum strategy for the attackers.

Figure 2 shows the proposed inferential model. The risk to the defender associated with a scenario is inferred from the aggregate consequence resulting from the attack and the likelihood that the attack is successful. As with the possibility tree, we use natural language expressions in the deduction of the inferential structure. The objective here is to deduce the factors determining risk and how they logically combine. The terminal nodes in this model are the input variables that are needed to evaluate the risk.

[b] When the entire tree of Fig. 1 is active, there are 2,702,151 unique attack scenarios.

TABLE I. Summary of Attack Scenarios

Scenario	Target	Radiation Source	Acquisition Source
S1	Confined Sports Arena	Gamma Source	Theft, insider diversion
S2	Confined Sports Arena	Gamma Source	Third party
S3	Downtown Area	Gamma Source	Theft, insider diversion
S4	Downtown Area	Gamma Source	Third party
S5	Downtown Area	Fission Products	Commercial Power Reactor
S6	Downtown Area	Improvised Nuclear Weapon	Theft, insider diversion

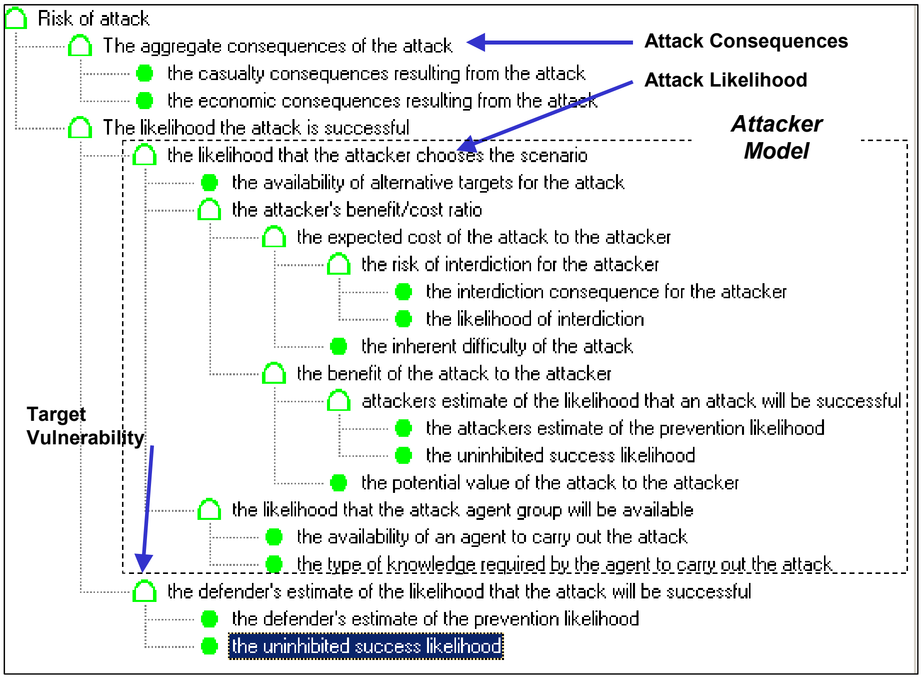
Risk of attack
The aggregate consequences of the attack
the casualty consequences resulting from the attack
the economic consequences resulting from the attack
The likelihood the attack is successful
the likelihood that the attacker chooses the scenario
the availability of alternative targets for the attack
the attacker's benefit/cost ratio
the expected cost of the attack to the attacker
the risk of interdiction for the attacker
the interdiction consequence for the attacker
the likelihood of interdiction
the inherent difficulty of the attack
the benefit of the attack to the attacker
attackers estimate of the likelihood that an attack will be successful
the attackers estimate of the prevention likelihood
the uninhibited success likelihood
the potential value of the attack to the attacker
the likelihood that the attack agent group will be available
the availability of an agent to carry out the attack
the type of knowledge required by the agent to carry out the attack
the defender's estimate of the likelihood that the attack will be successful
the defender's estimate of the prevention likelihood
the uninhibited success likelihood
Attack Consequences
Attack Likelihood
Attacker Model
Target Vulnerability

Fig. 2. Inference tree for attack risk.

Each AND gate in this tree is understood to represent an inferential step. We consider how this inference process is to be carried out in Section IV.

A simplified representation for the attacker’s decision process is shown as the Attacker Model in Fig. 2. The basic inferential model here is a cost/benefit analysis from the attacker’s perspective. Of course, the attacker’s actual decision model is inaccessible to us and we can only try here to approximate the analysis of a rational opponent. The attacker’s cost estimate depends upon the inherent difficulty of the attack as well as the risk of interdiction. Similarly in evaluating benefit, a rational attacker takes into account the perceived gain resulting from a successful attack and the likelihood that the attack is in fact successful. This latter node is inferred from the uninhibited success likelihood – the chances of success given no defense, and the attacker’s estimate that the defender can prevent the attack. The attacker must also assemble the appropriate agent group to perform the attack. In our simple model we picture some general pool of agents available to the attacker and shrink this pool depending upon the knowledge required to carry out the attack. The corresponding defender model here is a simple vulnerability estimate that depends again upon the uninhibited success likelihood and the likelihood that the attack can be prevented. The defender’s estimate of the chances of an attack succeeding in the absence of any preventive actions may well be different than what he believes to be the attackers estimate. The prevention likelihood will depend upon the preventive measures in place and their effectiveness, information that is directly available to the defender.

IV. APPROXIMATE REASONING EVALUATION FOR RISK

The inferential model identifies the basic factors and their relationships. At this point a decision must be made as to how the risk metric is to be measured using these factors. One approach is to represent the inferential process as a series of linked numerical functions. This is a natural choice when much of the input data is quantitative and where the data relationships can be easily cast in functional form. Consideration of the problem at hand suggests that these conditions are not met. Much of the data is qualitative in nature and the nature of the attacker/defender extensive game is primarily one of perception. An analysis methodology that is well suited to this set of inferential specifications is approximate reasoning (AR).[7,8] AR is intended to emulate the type of expert judgment used by subject matter experts. Variables in AR are linguistic – natural language expressions that can take on a set of discrete natural language values. Note that there is a one-to-one correspondence between the natural language used in the inference tree and these linguistic variables.

In AR, inferences with linguistic variables are made using formal logical implication defined in a series of linked, forward chaining rule bases. A detailed discussion of AR is beyond the scope of this paper and we offer here a short example to demonstrate the compatibility of AR with the type of expert-based inferential reasoning we wish to emulate. From Fig. 2 we conclude that risk, R can be inferred from the aggregate consequences C_a and the likelihood of a successful attack L_s,

$$C_a \wedge L_s \Rightarrow R \tag{1}$$

where $\Rightarrow$ is the implication operator. A linguistic variable takes on the values from its universe of discourse, the set of words used to describe it. Here we use R = {Very Low, Low, High, Very High}, C_a = {Negligible, Moderate, High, Catastrophic} and L_s = {Very Unlikely, Unlikely, Likely, Nearly Certain}. The rule base that implements the logical implication is shown in Table II. The shaded entry in the rule base corresponds to the implication “If the Aggregate Consequences are High and the Attack Success Likelihood is Unlikely Then the Risk is Low”. Each linguistic variable associated with Fig. 2 requires a universe of discourse and there is a rule base associated with each AND gate in this inferential model. The universes of discourse and the rule bases are chosen to be consonant with the judgments of the subject matter experts for the problem, in this case, intelligence officers, physical protection analysts, etc.

To perform evaluations with this AR model we treat each element in a universe of discourse as a fuzzy set. The ability to express ambiguous or hedged expressions of a linguistic variable, for example “The aggregate consequences are moderate to high, but somewhat closer to moderate” is encoded using the degree of membership vector, μ(Ca) = [Negligible:0, Moderate:.8, High:.2, Catastrophic:0]. The concept of likelihood appears a number of times in the risk model. Likelihood is a natural language expression associated with outcome uncertainty. We interpret a degree of membership for L_s , say $\mu(L_s)$ = {Very Unlikely: .2, Unlikely:.4, Likely:.3, Nearly Certain:0] as an expression of the possibility that the likelihood is described by each of these descriptors. Note that the sum of these memberships is not 1.0. That is, our expression of outcome uncertainty is imprecise and non-probabilistic. A natural language equivalent might be “the likelihood of a successful attack is in the range from very unlikely to likely”. Logical implication with fuzzy/ possibilistic inputs to a rule base is performed using the min-max operator.[9,10] For the two vectors given here, the corresponding risk vector is, μ(R) = [Very Low:.2, Low:.4, High:.2, Nearly Certain:0], which can be expressed variously as “the risk is low,” “the risk is very low to high,” etc. The application of linguistic variables, implication rule bases and fuzzy sets in AR provides the set of capabilities needed to evaluate risk for the nuclear terror problem.

V. ILLUSTRATIVE RESULTS

We examined the capabilities of the AR-based inference model by evaluating the six attack scenarios in Table I. This required providing degree of membership vectors for each of the terminal nodes in the inference model of Fig. 2. The values used were chosen to illustrate the characteristics of our game-based model and do not reflect actual intelligence or physical security data. Figure 3 shows the risk degree of membership vectors for the six scenarios. For each scenario the degree of membership in the risk fuzzy sets {Very Low, Low, High, Very High} are shown. Recall that S1 and S2 are attacks on an indoor sports arena that differ only by the acquisition mode. Scenarios S3 and S4 are the analogs for an attack on an outdoor downtown area. S5 is the scenario associated with a commercial power reactor and S6 is the improvised nuclear device; the target in these cases is a downtown area.

The distribution of the membership vectors arises from the assignments given to the input variables and evaluation of the chained implication rule bases. We may view these distributions as a representation of the ambiguity and outcome uncertainty associated with the scenario evaluations. It is often useful to present the results in terms of a centroid analogous to a mean or median in a probabilistic analysis.[c] These results are shown in Fig. 4. The numerical scale and the relative importance of the individual risk sets are chosen so that the results are consistent with the corresponding natural

[c] In an AR model this operation is referred to as defuzzification.

TABLE II. Rule Base for Inferring Risk from Aggregate Consequences and Likelihood of a Successful Attack

Agg.gate Conseq.	Catastrophic	*Low*	*High*	*Very High*	*Very High*
	High	*Low*	*Low*	*High*	*Very High*
	Moderate	*Very Low*	*Low*	*Low*	*High*
	Negligible	*Very Low*	*Very Low*	*Very Low*	*Low*
		Very Unlikely	Unlikely	Likely	Nearly Certain
		Likelihood of Successful Attack			

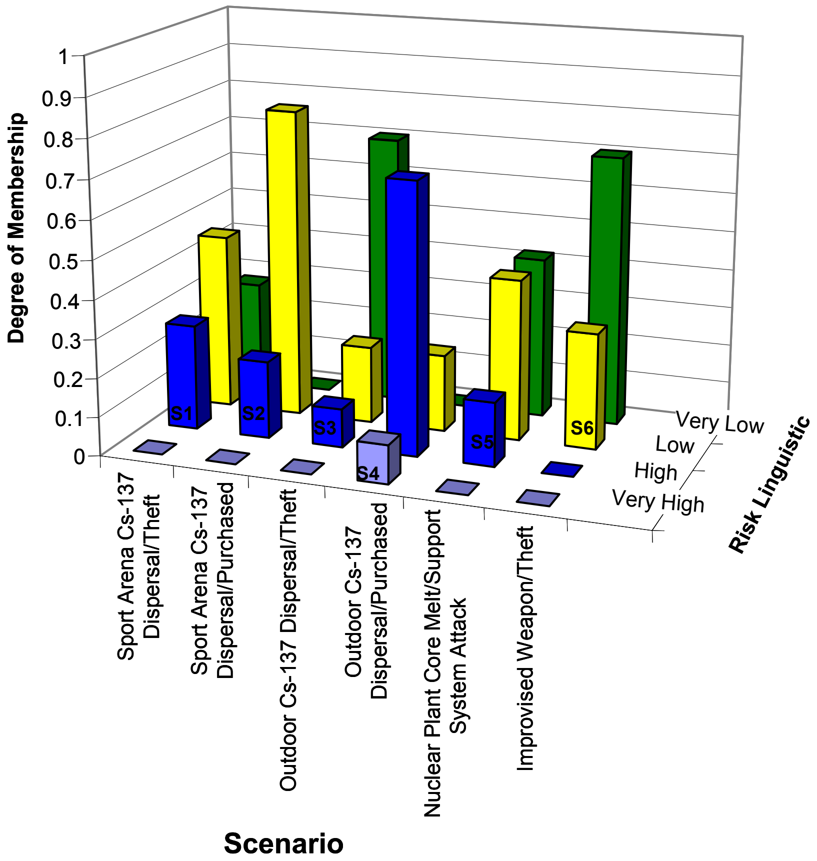

Fig. 3. Degree of membership vectors for illustrative attack scenarios.

language we use to describe risk. Figure 4 also shows a variation on the improvised nuclear device scenario. In S6* the scenario was modified so that the attacker is assumed to already have possession of the required special nuclear material in a location near the target. Although strictly notional here, the change in risk observed suggests the potential importance of nuclear safeguards in evaluating the risk of such attacks.

VI. DISCUSSION OF ANALYSIS

The application of LED to the problem of nuclear terrorism provides useful insights into the nature of homeland security. A possibility tree proved to be an efficient way to deduce a comprehensive set of attack scenarios. Although there are potentially an enormous number of unique scenarios, the use of termination and exclusion command options make it practical to select a subset of scenarios that is appropriate for detailed analysis. The natural language capability of LED TREE allows the attack scenarios to be expressed in a form that is accessible to the experts as well as decision makers. The details of a scenario can be developed in as much detail as desired. This allows for better identification, design and assessment of preventive and interdiction measures.

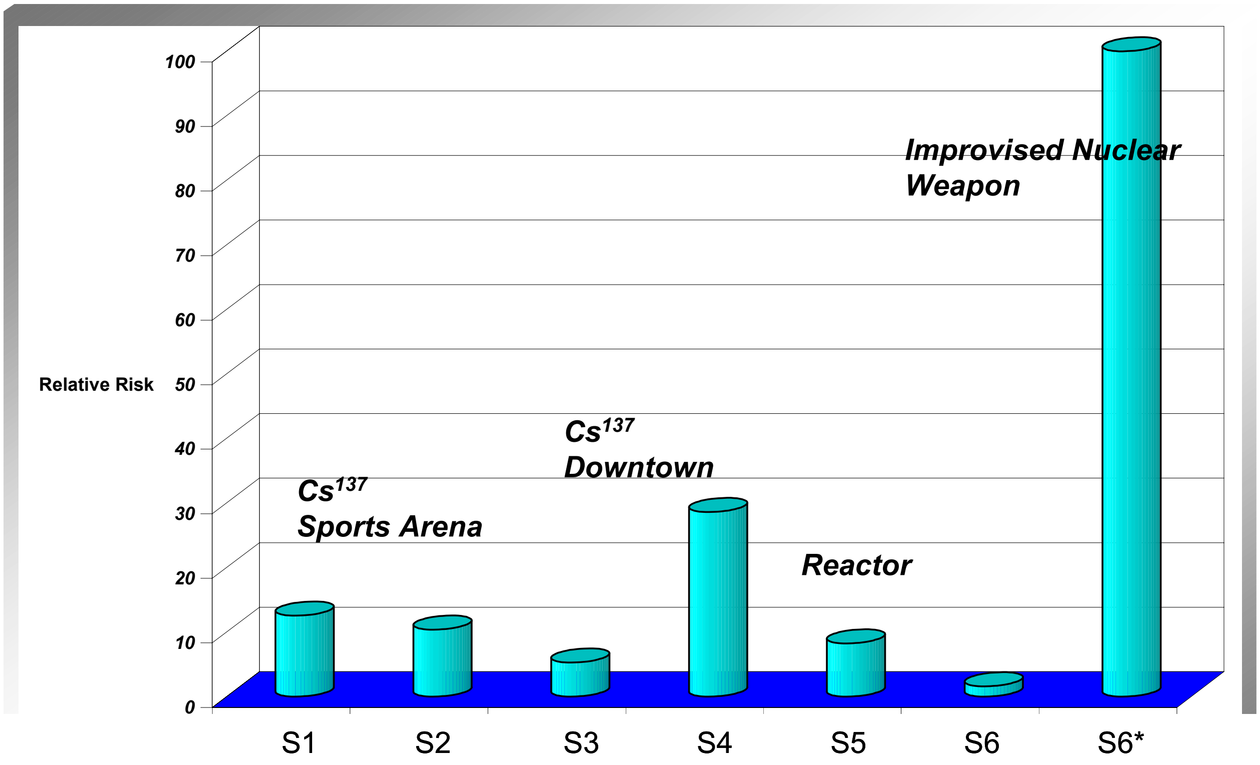

Fig. 4. Centroid risk values for illustrative problem.

The design of the inferential model presented here reflects our view that a game theoretic perspective is needed to make useful risk estimates. That is, it is insufficient to concentrate on the vulnerability of homeland targets to the exclusion of attacker motivation, intentions and capability. Efficient allocation of resources will not be possible unless realistic risk estimates are available. We believe that advanced game-based models that emulate attacker behavior will be important in achieving this goal. In our study we used an AR-based implementation to infer risk. AR works well with qualitative, perceptual data that is imprecise and ambiguous. Further, much of the inferential process attempts to reflect the intentions of players – terrorist groups, who are making their own sequence of expert judgments. It is relatively straightforward to represent judgments of this sort as a set of forward-chaining rule bases.

Other approaches such as Probabilistic Risk Analysis (PRA) have been suggested for estimating nuclear terror risk. LED software could be used to build the set of attack scenarios and to design an inferential model for a traditional PRA approach as well. We have discussed the relative merits of AR versus Bayesian approaches for complex, predominantly perception-based problems elsewhere.[11] AR is often “good enough” in the sense that an approximate rank ordering or triage of the alternatives is sufficient. For such problems high precision is neither practical nor useful. An AR approach meshes well with the types of intelligence information and analytical techniques that will play a prominent role in homeland security. Natural language methods such as AR are also compatible with data mining and other knowledge discovery techniques that are the prerequisites to a future capability to produce real time estimates of risk.

The results of our initial study are encouraging. An analysis of risk and asset allocation for a particular homeland security related problem is currently underway. Work on extending the capabilities of LED software to allow for more efficient model construction and evaluation is also underway. Finally we continue to pursue the objective of designing more realistic game-based inference models to improve risk estimation.

REFERENCES

1. K. B. CHRISTENSEN, T. F. BOTT, J. L. DARBY and S. W. EISENHAWER, **“**The Approximate Reasoning (AR)-Based Method for Information Loss Path Analysis (ILPA)**,”** *Proceeding of the 43rd Annual Meeting of the Institute for Nuclear Material Management*, (2002).
2. T. F. BOTT and S. W. EISENHAWER, “A Logic Model Approach to the Conceptual Design of a Scientific/Industrial Complex,” *Selected Topics on Aging Management, Reliability, Safety and License Renewal*, PVP- 44, 119, ASME, New York (2002).

3. S. W. EISENHAWER and T. F. BOTT, “Application of Approximate Reasoning to Safety Analysis,” *Proceedings 17th International System Safety Conference*, (1999).
4. S. W. EISENHAWER and T. F. BOTT, “Accident Reconstruction using Process Trees,” *Risk and Safety Assessment: Building Viable Solutions*, PVP- 320, ASME, New York (1995).
5. T. F. BOTT, S. W. EISENHAWER, J. KINGSON and B. P. KEY, “A New Graphical Tool for Building Logic-Gate Trees,” to appear *Proceedings of the 2003 ASME PVP Conference*, July 2003, Cleveland, Ohio.
6. M. J. OSBORNE and A. RUBENSTEIN, *A Course in Game Theory*, MIT Press, Cambridge (1994).
7. L. Zadeh, “A Theory of Approximate Reasoning,” *Machine Intelligence*, J. HAYES, D. MICHIE and L. MIKULICH, Eds., Halstead Press, New York (1976).
8. R. LOPEZ DE MANTARAS, *Approximate Reasoning Models*, Ellis Howrood, New York, (1990).
9. S. W. EISENHAWER, T. F. BOTT and R. E. SMITH, “An Approximate-Reasoning-Based Method for Screening High-Level-Waste Tanks for Flammable Gas,” *Nuclear Technology*, **130**, 351, (2000).
10. T. J. ROSS, *Fuzzy Logic with Engineering Applications*, McGraw-Hill, New York, (1995)
11. T. F. BOTT and S. W. EISENHAWER, “Evaluating Complex Systems when Numerical Information is Sparse,” *Proceedings of the World Automation Conference*, (2002).