Adversarial Risk Analysis for Enhancing Combat Simulation Models

Juho Roponen* and Ahti Salo

Systems Analysis Laboratory
Department of Mathematics and Systems Analysis
Aalto University School of Science
P.O. Box 11100, 00076 Aalto, FINLAND

*Corresponding author, email juho.roponen@gmail.com

Abstract

Adversarial risk analysis (ARA) builds on statistical risk analysis and game theory to help analyze decision situations which involve two or more intelligent opponents who make decisions under uncertainty. During the past few years, the ARA approach–which seeks to model the decision making processes of a rational opponent–has been applied extensively in areas such as counterterrorism and corporate competition. In the context of military combat modelling, however, ARA has not been used systematically, even if there have been attempts to predict the opponent's decisions based on wargaming, application of game theoretic equilibria, or solicitation of expert opinions. Against this background, we argue that combining ARA with military combat modelling holds significant promise for enhancing the capabilities of current combat modelling tools. Even if the ARA approach can be challenging to apply, it can be very informative because relevant assumptions about the resources, expectations and goals that guide the adversary's decisions must be clearly explicated. We identify some promising ways of combining ARA with combat modelling and present an illustrative example of how ARA can provide insights into a problem where the defender needs to estimate the utility gained from hiding its troop movements from the enemy.

Introduction

Adversarial risk analysis (ARA) combines statistical risk analysis and game theory to provide appropriate methods for analyzing decision making situations which involve two or more intelligent actors who make decisions with uncertain outcomes. Such situations are encountered, for example, in counter-terrorism and corporate competition (Rios Insua *et al.,* 2009).

Traditional statistical risk analysis was developed to assess and mitigate risks in contexts where the loss is governed by chance (or Nature), for instance in the management of complex technological systems like nuclear power plants and the design of insurance policies against natural disasters. Apart from risks caused by chance events, ARA seeks to capture risks caused by the self-interested and possibly malicious actions of intelligent actors, and consequently modelling the decision-making behavior of these actors is central to ARA. These kinds of decision models can be based, for example, on classical game theory (Myerson, 1991) or psychological considerations (Camerer, 2003).

Yet game theory is not ideal tool for describing and predicting human behavior. Minmax solutions–in which each actor seeks to minimize his expected losses across all the actions that are available to his opponents–can lead to sub-optimal solutions, because in reality opponents do not usually abide the minmax rationality principle. Minmax solutions are also often difficult to compute in real situations and necessitate strong assumptions what common knowledge the actors share (Kadane & Larkey, 1982 and Meng *et al.,* 2014). Moreover, the solutions can be overly too pessimistic, because the mitigation of the worst possible scenario (which may have an extremely low probability) will induce the actors to make choices that a human opponent would not realistically make.

ARA has many obvious uses in military organizations. Much of the recent ARA literature has focused on counterterrorism, and many of the proposed ARA approaches can be applied to support military decision making. Zhuang and Bier (2007), for example, apply game theory to devise strategies for the allocation of resources between the protection from an intentional attack, on one hand, and from natural disasters, on the other hand. ARA methods can be used to guide the allocation of resources between strategically important targets as well as the investment planning of military equipment and projects. Uses of ARA in finance and procurement are relevant, too, because military organizations acquire products and services from external contractors.

In this paper, we do not survey the broad relevance of the ARA literature in view of military applications. Rather, we discuss how ARA can be applied to enhance combat simulation or to serve as a complement to it. There are numerous combat models and simulators for calculating the outcomes of battles and the losses sustained by the units and weapons systems. Most of these simulators do not model human decision making except at the most basic level (Lappi, 2012, 14-20). Against this background, ARA holds potential for enhancing simulators by increasing the realism of the decision making processes that are embedded in the simulations, which, in turn, may help for instance in predicting the opponent's decisions.

One major advantage of ARA from the perspective of traditional combat modeling is the possibility to calculate the effects of military deception and its usefulness. Game theory has been applied to calculate the benefits of deceit before (Reese, 1980), but such applications are still rare. This is partly because the solutions of classical game theory presume that both sides have common knowledge about each other's goals and resources, which is not realistic when modeling deceit. ARA does not have this limitation. It can even be applied to calculate the usefulness of decoys and dummy systems, which makes it possible to estimate if they are worth the cost; this is very difficult for most combat simulation models to estimate.

Modeling adversarial risks

In this section, we briefly describe how a situation in which there are adversaries whose actions affect each other's risks can be modeled. Our analysis builds largely on the paper by Rios Insua *et al.* (2009) who give a comprehensive presentation of ARA. For a good overview on how the ARA approach compares to classical game theory, we refer to Banks *et al.* (2011).

Risk analysis

The simplest form of a non-adversarial risk management problem is a situation in which the decision maker (DM) chooses one alternative from the set of available decision alternatives whose costs are uncertain. These cost uncertainties may stem from the fact that the decision outcome is uncertain, or because the costs associated with a particular outcome are uncertain, or both. This problem is presented as an influence diagram in Figure 1.

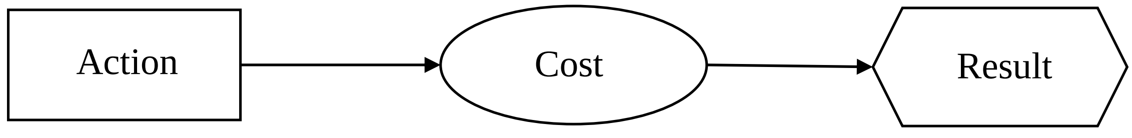

Figure 1. A simple influence diagram

An influence diagram is a directed acyclic graph with three kinds of nodes: rectangle shaped decision nodes, oval shaped uncertainty nodes, and hexagonal value nodes. Arrows pointing to value or uncertainty nodes indicate functional or probabilistic dependence, respectively. That means that the utility function at the value node depends on its immediately preceding nodes, and the probabilities associated with an uncertainty depend on the values of the immediately preceding nodes and are thus conditional on them. Arrows pointing into decision nodes indicate that the values of the nodes preceding the decision node are known at the time of the decision. (*cf.* Howard & Matheson, 2005)

The problem in Figure 1 represents a situation where the decision maker has to make a decision *a* from a set *A* of possible choices, represented by the rectangle. The cost *c* associated with this decision is uncertain and is modeled through density $\pi(c|a)$, represented by the oval node. The result is modeled by Von Neumann-Morgenstern utility function *u*(*c*). The decision maker seeks the decision that maximizes the expected utility

$$\psi = \max_{a\in A}\left[\psi(a) = \int u(c)\pi(c|a)\, dc\right]. \tag{1}$$

In practice, the costs of a particular action are complex and depend on the outcome. The costs often include fixed and random terms. For that reason, organizations seek to perform a risk assessment to better identify the disruptive events, and their probabilities and associated costs. Figure 2 shows the influence diagram that has been extended to account for the disruptive hazards identified by the risk assessment and the additional costs they may cause.

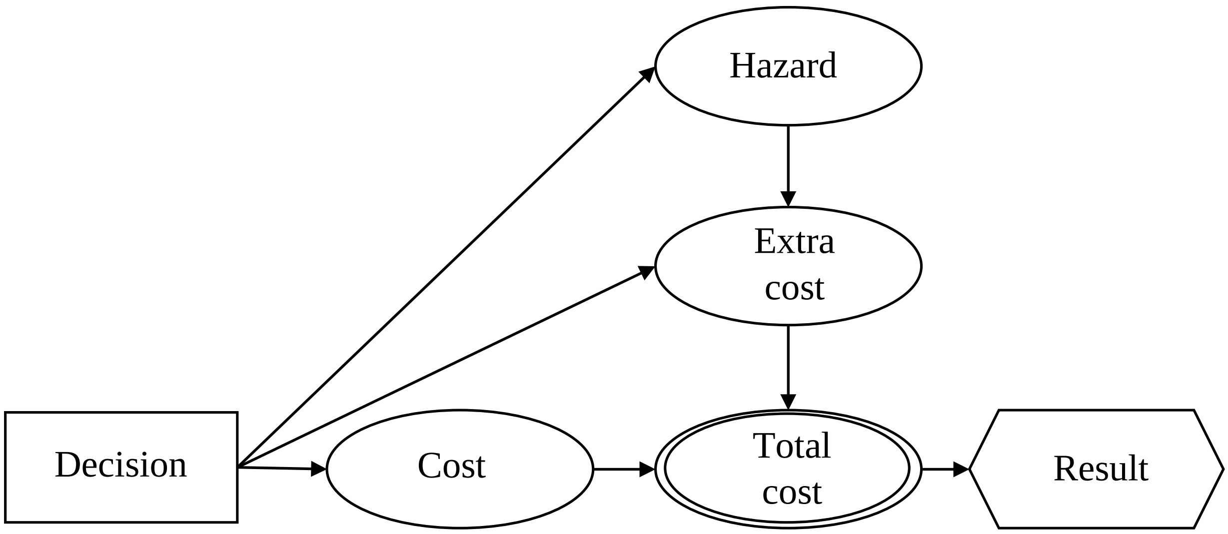

Figure 2. Influence diagram after risk assessment.

Adversarial risks

We now consider a situation in which there are two adversaries (Attacker and Defender) whose decisions affect the risks that each faces. Figure 3 extends the influence diagram to include the adversary in a symmetrical situation in which decisions of both parties affect the risks and costs that the other faces, and both seek to maximize their own expected utilities. Even though the roles are symmetric in this example, this does not have to be the case. An asymmetrical scenario could also be modeled with an asymmetrical influence diagram.

We denote the sets of possible actions of Attacker and Defender with *A* and *D* respectively. Their utility functions are $u_a(\cdot)$ and $u_d(\cdot)$. The sets containing their beliefs about different probabilities are

P_a and P_d. As can be seen in the influence diagram in Figure 3, one of the nodes, Hazard, is common to both sides. This can represent the possible complications arising from risks common to both sides, such as weather for example. The other cost nodes are not common, and represent the random costs for both parties and they could be very different.

The expected utilities for both the Attacker and the Defender depend upon the actions of both. Specifically, by extending on (1), we obtain the Attacker's expected utility for choosing action $a \in A$ when the Defender chooses action $d \in D$

$$\psi_A(a,d) = \int u_A(c)\pi_A(c|a,d)\,dc, \qquad (2)$$

where $\pi_A(c|a,d)\epsilon P_A$ represents the Attacker's beliefs about his costs corresponding to the decision pair (a,d). It is noteworthy that these beliefs do not necessarily have to match reality, because we are only modeling the decision the Attacker makes. The expected utility for the Defender is analogous.

This representation of ARA matches normal form games in which both players take simultaneous decisions. One could also build an influence diagram that applies to sequential games, such as Stackelberg games, in which the players make their move alternately. The ARA methodology can be applied to solve such games, too (*cf.* Banks *et al.*, 2011 and Rios & Rios Insua, 2012).

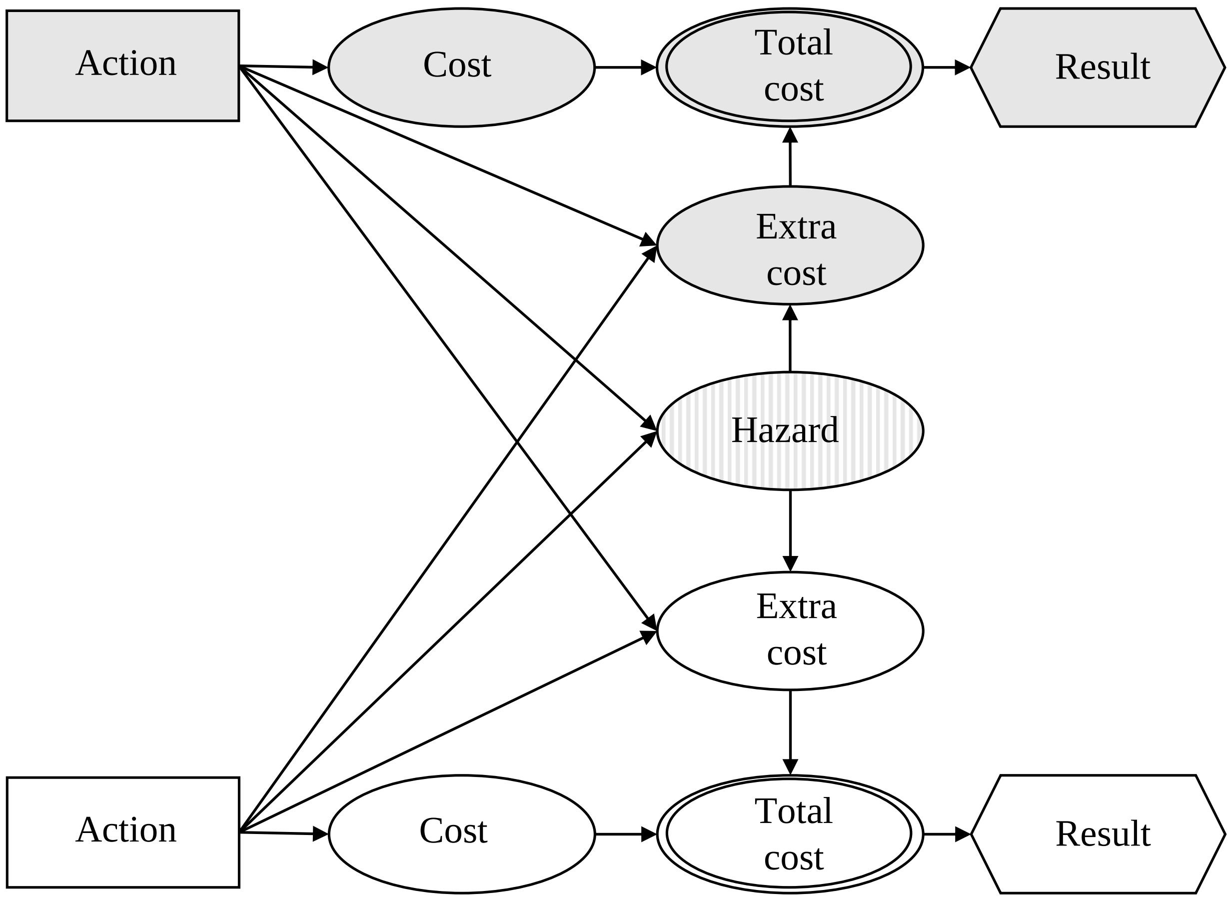

Figure 3. Influence diagram with an adversary

Bayesian framework for ARA

A problem like the one presented in Figure 3 can be readily solved using classical game theory if we assume that the costs and utility functions of both players are common knowledge. If the players do not possess correct and accurate information about the costs, resources, and goals of the adversary, (which is often the case in reality), the Nash equilibrium solution does not exist.

ARA solves this problem by using a Bayesian strategy to express uncertainty about the adversary's decision. If we examine the problem from the Attacker's point of view, the uncertainty means that the Defender's decision is a random variable as presented in Figure 4. To solve this problem, the Attacker needs more than just $\pi_A(c|a,d) \epsilon P_A$ and $u_A(c)$. Specifically, he also needs $p_A(d)$, which is the probability that the Defender chooses defense d as estimated by the Attacker. To find that, he is assumed to use mirroring to form an estimate of both the Defender's utility function $u_D(c)$ and the Defender's costs $\pi_D(c|a,d)$. That means that the Attacker assumes the Defender is acting rationally and is using a similar strategy to predict the actions of the Attacker.

If the Attacker tried to estimate the Defender's utility function and cost function by assuming that the Defender is doing the exact same thing that he is doing, the Attacker would need to think what the Defender thinks he thinks. To avoid infinite regress, the chain is usually cut there and the Attacker just forms an educated guess about the Defender's thoughts about the Attacker's estimated utilities and costs. Obviously the thinking could be taken even further, but it usually not a realistic way to resolve the problem.

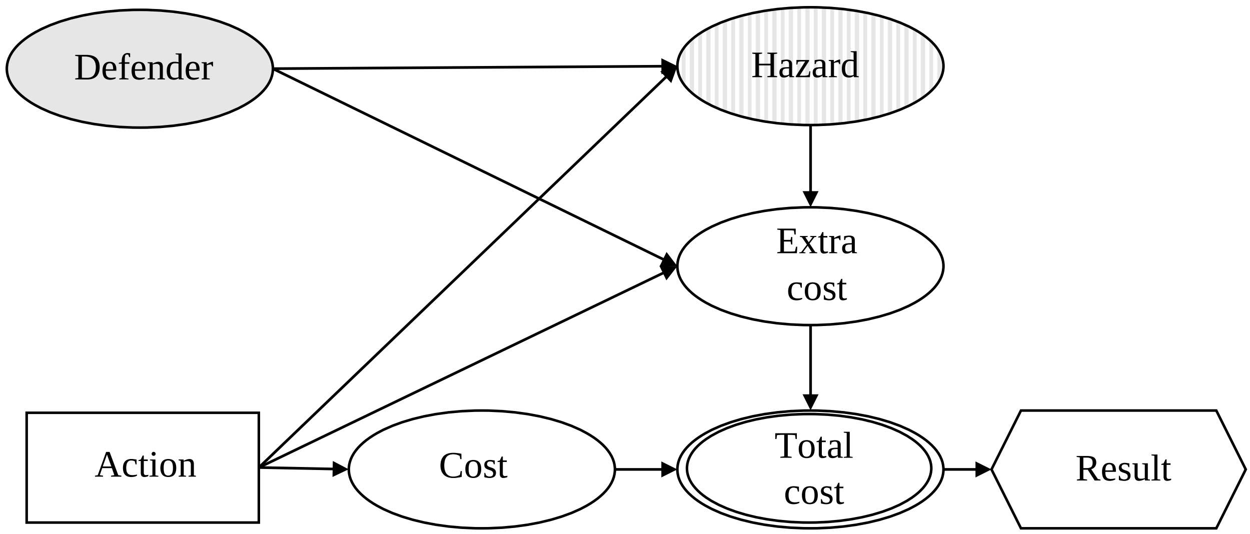

Figure 4: Influence diagram from the Attacker's point of view.

Alternative approaches for modeling adversary's decision making

The ARA methodology is quite similar to Bayesian level-k thinking. The approach to the modeling of opponents thoughts as outlined here resembles most closely level-2 thinking. Rothschild *et al.* (2012) have taken the approach further and applied actual level-k thinking to the ARA approach. Their methodology is not without drawbacks, because the level-k approach requires some additional assumptions and the problems become intractable even more rapidly due to their increasing complexity. Possibly the greatest advantage gained from the level-k thinking approach is the ability to easily perceive how the level of adversary's thinking affects the optimal decision.

Caswell *et al.* (2011) have presented a model that evaluates the decision process using a Bayesian network with an embedded semi-Markov decision process. Compared to the ARA approach their model can be used to present the adversary's decision process with greater accuracy. However, as with any decision analysis model the results are only as good as inputs, and a detailed description of the adversary's thought process would also require extremely detailed information about the adversary's resources, values and goals.

Zuckerman *et al.* (2012) have taken a more distinct approach by modeling adversarial activity using Beliefs-Desires-Intentions (BDI) based model, which are commonly used to describe teamwork and cooperation. This approach also has the ability to model the adversary as a more nuanced rational agent instead of an omniscient utility maximizer. The model is still somewhat unrefined as it is only applicable in zero-sum games like environments with easily decomposable goals, but it is a good starting point.

Applying ARA to military combat modeling

In this section, we discuss possibilities of applying ARA to military combat modeling and modeling processes. These ideas are still mostly untested, and they are presented as suggestions for worthwhile topics for future research.

Distribution of resources

A significant proportion of ARA literature is focused on preventing terrorist threats and, more specifically, on how limited resources should be best allocated to combat these threats (*cf.* Pat-Cornell & Guikema 2002, Kardes & Hall 2005, Zhuang & Bier 2007, Golany *et al.* 2009, and Kroshl *et al.* 2015). Resource allocation is a problem that military faces always, not just when combating terrorism. The methods developed for allocating resources against terrorist threats are easy to apply to combat environment, because one of the biggest challenges faced when using these methods is assessing the effects of each of the adversary's possible decisions. Combat modeling already has the tools for estimating the results of combat when the adversary has committed to a specific strategy.

This is probably the easiest way to benefit from ARA in the context of combat modeling. Studying the ARA research related to counter terrorism and applying ARA methods to combat situations by calculating the risks related to each possible adversary's decision using combat modeling tools offers potentially significant results without any new model development.

When combined with data farming ARA approach can also be used to calculate how limited resources such as troops should be divided across different battlefields. Kovenock & Roberson (2010) for example have examined using game theory how resources should be allocated for multiple battlefields when the objectives and available resources for both sides are known. ARA methodology can be used to perform similar analysis with fewer assumptions.

Modeling decision making

Simple adversarial intent models have been used in professional wargaming to simulate intelligent forces (Santos & Zhao 2006). Despite that, currently many combat models do not include algorithms that are capable of modeling the human thought process behind tactical or higher level decisions (Washburn & Kress, 2009, 111-130). Depending on the model, practically all higher level decisions (above the levels of soldiers or platoons) are made by the operator. As a result, the time required to create a scenario is usually significantly longer than the time required to calculate the results (Lappi, 2012).

The fact that the operator user defines the actions taken by the troops in the modeled scenario is not always a hindrance. It also offers benefits, including the ability to fine tune the scenario to match events of a historical battle (Lappi *et al.*, 2014) or to easily examine and change the actions taken by the troops. Relying on the operator in making all the significant decisions can, however, become problematic when a really large number of scenarios is required, which is often the case when data farming more complex cases. (Lappi *et al.*, 2015) Varying the scenario is done by modifying numerical parameters such as troop strengths or number of available weapon systems, but that approach is somewhat limited in its applications. In longer and more complex scenarios small changes that affect outcome of the battle in its early stages can for example cause the assumptions made about the troop movements in the later stages become invalid.

Including a model that would allow the units inside the simulation make simple tactically sensible decisions would widen the range of problems that can be analyzed using data farming. The ARA methodology could be used as a basis for such decision making algorithm. Here, ARA has advantages over using game theory to calculate a Nash equilibrium, because the ARA framework makes it possible to better account for uncertainties and possible misinformation that are present on the battlefield. For a Nash equilibrium to exist both players need to have more information about the objectives and resources of the opponent than is often realistic in a combat situation.

However, there are some limitations to applying ARA to decision modeling. Because the method is so calculation intensive, it is not ideal for modeling low level or continuous decision making. The most effective way to apply ARA would be to limit the choice to between a few possible strategies. In contrast, considering very frequent decisions with a large number of distinct alternatives would lead to intractable models. Simplifying the problem by reducing the choices to move sequences instead of trying to calculate all possible actions has been long used in constructing AI systems for games such as Go and Chess and it has been recently applied even in video game AI development (Churchill *et al.* 2012).

Simulating larger chains of events

The ARA methodology can also be applied to aid in modeling military operations that are too large to simulate as a single scenario. The scale can become an issue if the number of units involved is too large, or the operation takes place over such a long timeframe that the number of possible paths based on the events becomes excessive. Kangas and Lappi (2006) present how methods of probabilistic risk analysis can be used in conjunction with stochastic combat modeling to analyze longer chains of events. The ARA approach could be used to build on those results and take the analysis one step further. In addition to predicting the success chances of larger operations, it would also be possible to predict those of the adversary’s choices that can affect the path of events.

Practically any combat model can be used with ARA methodology on condition that the probabilities for each side winning the battle as well as the expected losses on both sides can be calculated. This includes practically all stochastic combat models and even some of the deterministic ones. The selection of the combat model must fit the problem at hand. Sometimes the best choice is a platform level Monte Carlo simulation, and sometimes it is a high level attrition model like FATHM (Brown & Washburn, 2000).

In some cases, it can be possible to use ARA to model these longer chains of events without having to rely on an actual stochastic combat modeling software like Sandis (cf. Kangas and Lappi, 2006). There are also alternative, lighter stochastic computational models that can be used to predict the outcome of a duel between two platoon sized forces (Lappi *et al.,* 2012; Åkesson, 2012; Roponen, 2013). These models can be used to significantly cut down the time for calculating all the success probabilities and the expected losses in different stages of the chain. There are also additional time savings from not having to create a complete model scenario, which, as noted earlier, is a time consuming process. The use of the lighter duel simulation methods could even be automatized to a certain degree, because they require far fewer input parameters.

Modeling the effectiveness of military deceit

Using deceit to gain upper hand against an adversary is an absolutely integral part of military tactics and strategy. Still, the effects of deceit are very difficult to predict and simulate with existing operational analysis and combat modeling software. Because the effects cannot be readily reduced to mathematical formulas, modeling the effects of deceit relies usually on expert opinions. In the context of combat modeling, this usually means that the required expert opinions are provided by the software operator.

A common alternative is to use wargames to model the uncertainties associated with human decision making, but this approach also has some problems as even wargames are forced to ignore certain aspects of reality. Questions of solvability do not arise in wargames because optimal tactics are not sought. Because wars are fought by humans, humans are also used to model the decision process. The first problem that stems from this approach is, that the player can make decisions in a game that he would not make in a real world as long as it produces good results in the simulation. For example casualties might not carry the same weight in a simulated environment. The second problem is that wargames often capture typical decision making instead of optimal by having players only play a very small number of games. Lack of repetition overstates the effectiveness of new weapon systems, because the opponent does not have time to learn and adapt his tactics to counter those. The lack of repetition is deliberate to some extent as it is feared that the players would learn to use the artificialties of the wargame to their advantage instead of developing better military strategies. Another reason for the lack of repetitions is that wargaming is time consuming and expensive. (Washburn & Kress, 2009, 111-130)

The ARA approach could be used to assess the effects that deceit tactics could have on the decision making of the adversary. Specifically, the ability to model the effects of the adversary's altered perceptions would be a very useful complement to the elicitation of expert opinions. Mathematical equations are, after all, immune to effects of optimistic thinking.

Examples of situations that could be modeled with ARA relatively easily include cases in which the adversary is deliberately misinformed about the strength of the opposing forces. This can be achieved for example by hiding troop movements and employing dummy units or decoys. ARA can be then used to estimate the effect of the deceit on adversary's decision making and whether that effect is beneficial or not. An example of such estimation process is given in the next section.

Supporting decision making

Arguably the most important reason for military combat modeling is its use for supporting strategic, tactical or technical decision making process (Tolk, 2012, 55-78). However, it is not an easy task to translate the results of combat models into actual decisions or recommendations (Davis & Blumenthal, 1991).

The ARA framework could be applied to translate the data, produced by the mathematical models, to answer more concrete questions such as what will happen if we do not allocate more troops to a specific airfield, or where is the enemy likely to attack if we do X? As the answers to these questions depend on the decisions made by the opponent they are mathematically and computationally very difficult, but it is possible to answer them using the ARA methodology. (What is prohibitively hard to calculate today may not be so five years from now, thanks to the rapid increase in computing power.)

One way of using ARA to translate simulation results to an easily applicable form is to perform a portfolio analysis of all relevant strategies. Similar methods have already been used in assessing cost-efficiencies of different weapon system combinations (Kangaspunta *et al.*, 2012). In the same vein, ARA could be used to predict the most likely responses of the adversary and to calculate the expected utilities of each strategy under different conditions. The applicability of this approach depends on the ability to streamline and automate the process so that informative results concerning strategies can be provided while saving time in comparison with manual analyses of combat simulations.

Example of applying ARA to a tactical problem

To demonstrate how ARA can be applied in practice, we use it here to solve a relatively simple tactical problem with the aim of giving a general idea of the methodology which, we believe, can be applied to much more complex situations.

The problem

Consider the following situation in which there are two adversaries: the Defender and the Attacker. The Defender is moving in more troops to protect a valuable target, and the Attacker has to decide whether he will try and capture the valuable target or use is his troops somewhere else.

To make the problem easier to understand and calculate, we have reduced the decision of the Defender and the Attacker to simple binary choices. The Defender decides whether or not to hide the movements of his reinforcements from the Attacker. The Attacker will decide, after observing the perceived strength of the defender, whether or not he will attack the target. If the Defender decides to hide his reinforcements, the Attacker will not know about them and will therefore decide whether or not to attack based on incomplete information. We solve this problem from the point of view of the Defender. Figure 5 shows the influence diagram from the Defender’s point of view.

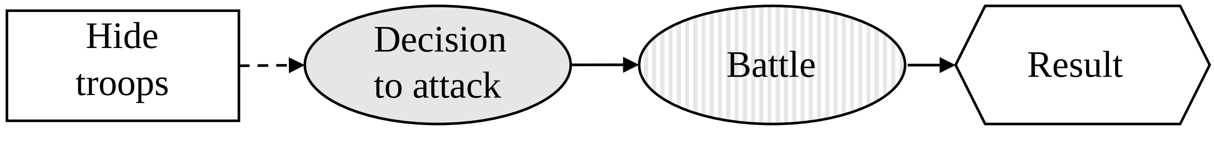

Figure 5. Influence diagram of the example case from the point of view of the Defender.

We denote the options the Defender and the Attacker as $D = \{0,1\}$ and $A = \{0,1\}$ respectively. The only uncertainty in this case is the outcome of the battle S (say, success or failure for the Defender to keep hold of the target). The utility functions over the costs are $u_D(c_D, c_A)$ and $u_A(c_D, c_A)$, with costs dependent on the actions of the Attacker.

In order to solve this problem, it is necessary to assess the probabilities over the costs, conditional on (a, d, S); and about S, conditional on (a, d). In this case, the Attacker and the Defender have different assessments: for example, for success, these are $p_D(S = 1|a, d)$ and $p_A(S = 1|a, d)$, respectively. It is likely that the Attacker's assessment of the success of the assault differs from that of the Defender, because the Attacker is not aware of any choices being made by the Defender (see Figure 6). The expected utility for the Attacker, resulting from (a, d) is

$$\begin{aligned}\psi_A(a, d) = p_A(S = 0|a, d) \sum_{c_A} \sum_{c_D} [u_A(c_A, c_D)\pi_A(c_A, c_D|a, d, S = 0)] \\ + p_A(S = 1|a, d) \sum_{c_A} \sum_{c_D} [u_A(c_A, c_D)\pi_A(c_A, c_D|a, d, S = 1)].\end{aligned} \tag{3}$$

The Defender's expected utility is similar.

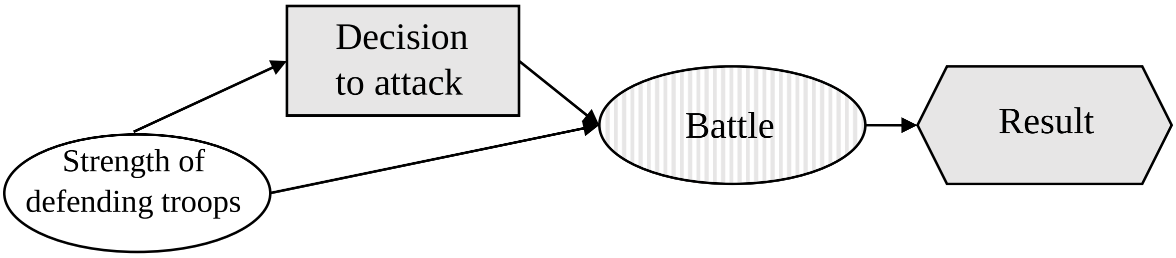

Figure 6. Influence diagram of the example case from the point of view of the Attacker.

We now solve the game from the Defender's point of view. The Defender has 15 men defending the target and has 15 more men coming in as reinforcements. He has the option of hiding the presence of the reinforcements from the Attacker. The Defender estimates that the Attacker has at least 20 men but no more than 35, and he thinks that the most likely number is 30, so he fits a triangular distribution as seen in Figure 7.

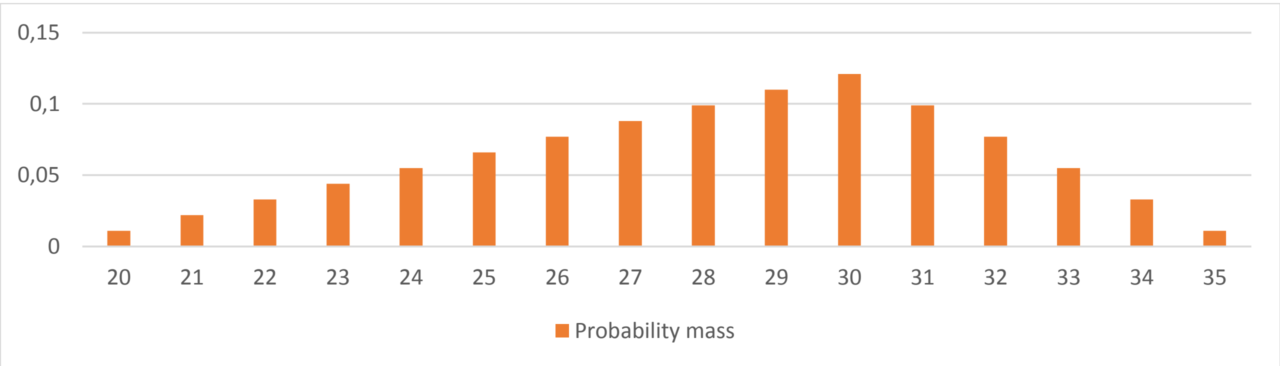

Figure 7: Defender's estimation of Attacker's strength as a probability mass distribution.

Using the strength estimates of both forces, the Defender can use, for example, a stochastic combat model to calculate p_D and π_D. Specifically, the Defender assesses that the utility gained from the situation follows the function

$$u_D(S, a, c_A, c_D) = 35S + 5a + 0.1c_A - c_D, \tag{4}$$

where $S = 1$ corresponds to the situation where the Defender manages to hold the target area, $a = 1$ corresponds to the situation where the Attacker decides to attack, instead of using the troops elsewhere, and c_A and c_D are the Attacker's and the Defender's losses respectively.

It is not enough to solve the problem for the Defender to know p_D, π_D and u_D. To calculate the expected utility from decision d, he first needs to estimate $p_D(a|d)$. Towards this end, the Defender needs to solve the problem from the viewpoint of the Attacker. He assumes that the Attacker is also an expected utility maximizer. The problem is presented from the Attacker's point of view in Figure 6.

The Defender estimates that the Attacker thinks the Defender has 13 to 17 men (with all values equally probable), and will not find out about the reinforcements if the Defender decides to hide their movement. If the Defender decides not to hide the reinforcements, he estimates the Attacker will think the Defender has 28 to 32 men (again, all values in this range equally probable). Using those strengths for his estimates he can use the same stochastic combat model used to solve p_D and π_D to calculate p_A and π_A.

The Defender estimates that the Attacker's utility function is similar to his own. However, the Defender does not know for sure how valuable the target is to the Attacker and what is the loss of opportunity that the Attacker suffers from not being able to use the troops elsewhere. He models this uncertainty by adjusting the weights of successful attack S and the decision to commit troops to the effort a. Thus, he estimates that the Attacker's utility function is

$$u_A(S, a, c_A, c_D) = -(35 + U_1)S - (5 + U_2)a + 0.1c_D - c_A, \tag{5}$$

where U_1 and U_2 are assumed uniform on $[-5, 5]$ to make calculations simpler.

Let us look at solving the problem step by step. To solve the problem, the Defender will:

1. Calculate the success probabilities and expected losses for both sides for all the possible combinations of strengths of both sides as perceived by the attacker.
2. Calculate the Attacker's expected utilities ψ_A for attacking and not attacking for all possible strengths of the Attacker's force taking into account the uncertainties with u_A.
3. Compare the expected utilities to get an estimate for the probability of an attack for each possible strength of the Attacker as seen in Figure 8.
4. Consider the probability of an attack with a specific strength of the attacker (Figure 8) and the probability for each of those strengths (Figure 7) to calculate $p_D(a|d)$.
5. Calculate ψ_D for all possible values of a.
6. Use $p_D(a|d)$ to determine the decision d which maxizes his expected utility.

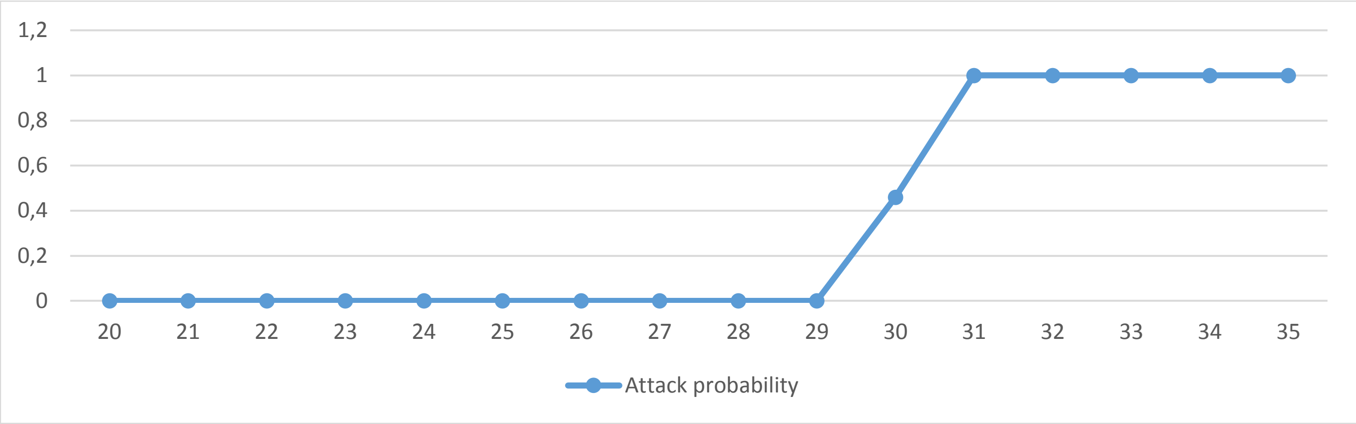

Figure 8: Probability of an attack as a function of the Attacker's strength, when the Defender has not hidden his reinforcements.

We used the approximative method in Roponen (2013) to simulate a duel between two forces in order to calculate p_D, π_D, p_A and π_A, because this program code for this method was available and provided the results quickly and with sufficient accuracy. This program code examined all the possible strengths of both sides, calculated the expected utilities and determined the attack probabilities $p_D(a|d)$. The attack probability when the Defender hides the reinforcements $p_D(a|d = 1) = 1$, and when he hides the reinforcements $p_D(a|d = 0) \approx 0.33$. Then the expected utilities of the Defender were calculated from p_D and π_D as

$$\psi_D(a, d = 1) \approx 13.44, \text{ and } \psi_D(a, d = 0) \approx 27.88. \qquad (6)$$

Thus the Defender decides that hiding the reinforcements is not in his best interests.

The result may seem counterintuitive; yet it does make sense when one realizes that (i) the safety of the target is of greater value to the Defender more than the losses sustained by the Attacker and (ii) the Defender does not necessarily outnumber the Attacker even with the reinforcements at hand.

Discussion

Adversarial risk analysis (ARA) is still a relatively new research area which is becoming more prominent in the context of counter-terrorism and corporate competition. As we have argued in this paper, ARA has much to offer for military combat modeling, not least because ARA is able to combine the statistical approach of risk analysis–which is already widely employed in combat modeling–with a game-theoretical perspective that helps predict the actions of one's opponents.

In particular, we have outlined possible uses for the ARA approach in the context of military combat modeling, operations analysis, and decision making in general. The ideas in this paper are still tentative and call for more research and development before they can be fully implemented. Real problems are extremely complex and solving them using the ARA methodology can be challenging. Actual battles involve thousands of decisions, and there are major uncertainties about the goals and resources. Yet, by selecting those decision contexts in which ARA is particularly appropriate will make it possible to benefit from the major opportunities that ARA has to offer.

We also presented a relatively simple example in which ARA was combined with stochastic combat modeling. In this example, most calculations for solving the ARA part of the model were relatively straightforward and could be readily implemented into program code (there are numerous tools for calculating the outcomes of battles; see, Kangas, 2005). This suggests that it is possible to develop software tools for problems that are more complex than in our example by building well-founded models of how the opponent's utilities affect his decision making. Indeed, there is much potential in using the ARA approach to tackle realistic problems in the context of stochastic combat modelling. This, we believe, would help push the boundaries of ARA research in this important application area.

References

Banks, D., Petralia, F., & Wang, S. (2011). Adversarial risk analysis: Borel games. *Applied Stochastic Models in Business and Industry 27*(2), pp. 72-86.

Brown, G. & Washburn, A.R. (2000 rev. 2004). *The fast theater model (FATHM)*, Project Report, (NPS-OR-01-002-PR). Naval Postgraduate School, Monterey, CA.

Camerer, C. (2003). *Behavioral Game Ttheory: Experiments in Strategic Interaction*. Princeton University Press, New Jersey.

Caswell, D. J., Howard, R. A., & Paté-Cornell, M. E. (2011). Analysis of national strategies to counter a country's nuclear weapons program. *Decision Analysis*, *8*(1), 30-45.

Churchill, D., Saffidine, A., & Buro, M. (2012). Fast Heuristic Search for RTS Game Combat Scenarios. In *Proceedings of the Eighth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, AIIDE.*

Davis, P. K., & Blumenthal, D. (1991). *The base of sand problem: A white paper on the state of military combat modeling* (No. RAND/N-3148-OSD/DARPA). Defense Advanced Research Projects Agency, Arlington VA.

Golany, B., Kaplan, E. H., Marmur, A., & Rothblum, U. G. (2009). Nature plays with dice–terrorists do not: Allocating resources to counter strategic versus probabilistic risks. *European Journal of Operational Research*, *192*(1), 198-208.

Howard, R. A. & Matheson, J. E. (2005). Influence diagrams. *Decision Analysis*, *2*(3), pp. 127-143.

Kadane, J. B. & Larkey, P. D. (1982). Subjective probability and the theory of games. *Management Science*, *28*(2), pp. 113-120.

Kangas, L. (2005). Taistelun stokastinen mallinnus. Master's Thesis, Helsinki University of Technology. http://sal.aalto.fi/publications/pdf-files/tkan05.pdf Accessed: 2014-05-07

Kangas, L. & Lappi, E. (2006) Probabilistic risk analysis in combat modeling. In: Hämäläinen, J. (ed.) *Lanchester and Beyond. A Workshop on Operational Analysis Methodology*. PVTT Publications 11.

Kangaspunta, J., Liesiö, J., & Salo, A. (2012). Cost-efficiency analysis of weapon system portfolios. *European Journal of Operational Research*, *223*(1), pp. 264-275.

Kardes, E., & Hall, R. (2005). Survey of literature on strategic decision making in the presence of adversaries. Unpublished report.

Kovenock, D., Roberson, B. (2010). Conflicts with multiple battlefields, *CESifo working paper: Empirical and Theoretical Methods, No. 3165*

Kroshl, W. M., Sarkani, S., & Mazzuchi, T. A. (2015). Efficient allocation of resources for defense of spatially distributed networks using agent-based simulation. *Risk Analysis* (to appear).

Lappi, E. (2012). Computational methods for tactical simulations. *Julkaisusarja 1. N:o 1/2012*. Doctoral Thesis, National Defence University, Finland.

Lappi, E., Pakkanen, M., & Åkesson, B. (2012). An approximative method of simulating a duel. In: *Proceedings of the Winter Simulation Conference*, *WSC '12*, pp. 208:1-208:10'

Lappi, E., Pentti, J., Åkesson, B., Roponen, J., Valtonen, J., Koskinen, J., Burhan, U., Sivertun, Å., & Hämäläinen, J. (2015). Team 4: Data farm. Manuscript.

Lappi, E., Urek, B., Åkesson, B., Arpiainen, J., Roponen, J., Jokinen, K., & Lappi, R. (2014). Comparing simulated results and actual battle events from 1944 – A case study using Sandis software. *Tiede ja Ase No. 72*, pp. 111-125.

Meng, S., Wiens, M., & Schultmann, F. (2014). A game-theoretic approach to assess adversarial risks. In: Brebbia, C.A. (ed.), *Risk Analysis IX*, WIT Press, Southampton, UK, pp. 141-152.

Myerson, R. B. (1991). *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, MA.

Paté-Cornell, E., & Guikema, S. (2002). Probabilistic modeling of terrorist threats: A systems analysis approach to setting priorities among countermeasures. *Military Operations Research*, *7*(4), pp. 5-23.

Reese, W. (1980). Deception in a game theoretic framework. In: Daniel, D. C., Herbig, K. L., Reese, W., Heuer, R. J., & Sarbin, T. R. (1980). *Multidisciplinary Perspectives on Military Deception* (No. NPS-56-80-012A). Naval Postgraduate School, Monterey, CA.

Rios Insua, D., Rios, J., & Banks, D. (2009). Adversarial risk analysis. *Journal of the American Statistical Association*, *104*(486), pp. 841-854.

Rios, J., & Rios Insua, D. (2012). Adversarial risk analysis for counterterrorism modeling. *Risk Analysis*, *32*(5), pp. 894-915.

Roponen, J. (2013). Kaksintaistelun approksimatiivinen mallintaminen. Bachelor's Thesis, Aalto University. http://sal.aalto.fi/publications/pdf-files/trop13_public.pdf Accessed: 2014-05-07

Rothschild, C., McLay, L., & Guikema, S. (2012). Adversarial risk analysis with incomplete information: A level-k approach. *Risk Analysis*, *32*(7), pp. 1219-1231.

Santos Jr, E., & Zhao, Q. (2006). Adversarial models for opponent intent inferencing. In: Kott A., & McEneaney, W. (2006) *Adversarial Reasoning: Computational Approaches to Reading the Opponents Mind.* Chapman & Hall/CRC, Boca Raton, pp. 1-22.

Tolk, A. (2012). *Engineering Principles of Combat Modeling and Distributed Simulation*. Wiley. USA.

Washburn, A. R., & Kress, M. (2009). *Combat Modeling*. Springer. Heidelberg.

Zhuang, J., & Bier, V. M. (2007). Balancing terrorism and natural disasters-defensive strategy with endogenous attacker effort. *Operations Research, 55*(5), pp. 976-991.

Zuckerman, I., Kraus, S., & Rosenschein, J. S. (2012). The adversarial activity model for bounded rational agents. *Autonomous Agents and Multi-Agent Systems*, *24*(3), 374-409.

Åkesson, A. (2012). Automatic calculation of win probabilities and conditional strength distribution of units in stochastic simulation model. Master's Thesis, Åbo Akademi University.