Optimisation, Games, Adaptation: Three Perspectives on Operations Research for Counter-IED

Dr. Anthony H. Dekker
Joint Operations Division
Defence Science and Technology Organisation (DSTO)
Canberra, Australia
dekker@ACM.org

Abstract. This paper explores Operations Research issues in the response to Improvised Explosive Devices (IEDs), using the concept of a "fitness landscape." In particular, we examine optimisation approaches that assume a fixed fitness landscape for Blue actions; game-theoretic approaches where fitness is associated with the combination of Red and Blue actions; and approaches that assume fitness landscapes are constantly changing as a result of Red and Blue adaptivity. We discuss the strengths and weaknesses of these approaches with respect to an illustrative simulation model, and present experiments suggesting that genetic algorithms are a promising mechanism for exploring adaptivity in such simulation models.

1. INTRODUCTION: IEDS

Improvised Explosive Devices (IEDs) have had a significant impact on recent military operations by a number of countries (JIEDDO, 2008; Zorpette, 2008a, 2008b). For example, Figure 1 shows the impact of an IED blast.* As well as causing deaths, IEDs are also responsible for many injuries. They remain an issue in several ongoing conflicts.

Insurgents may construct and place IEDs in a number of different ways, and may trigger them either directly (by using a radio or command wire), or indirectly (when the victim activates a pressure plate or infrared sensor).

The ongoing IED problem leads to a need for effective Counter-IED Operations Research (OR). In this paper we use a simple agent-based simulation, written in Java, to illustrate the strengths and weaknesses of OR approaches based on **optimisation**, on **game theory**, and on **adaptation**.

Figure 1: IED aftermath: a Stryker vehicle overturned by a buried IED blast (photo from www.army.mil)

* All data in this paper is taken from open sources, such as books, journals and the Internet.

2. A SIMPLE SIMULATION MODEL

As an illustration of some of the Operations Research (OR) issues involved in analysing Counter-IED operations, we have constructed the very simple agent-based simulation shown in Figure 2.

This simulation is not intended to accurately represent real-world IEDs or Counter-IED activities, but rather to illuminate the Operations Research issues they raise.

Blue agents must traverse a 20×20 grid from left to right. The grid contains four kinds of terrain: a gently curving **road**, a meandering **path**, a large expanse of **sand**, and randomly-placed **rocky** areas.

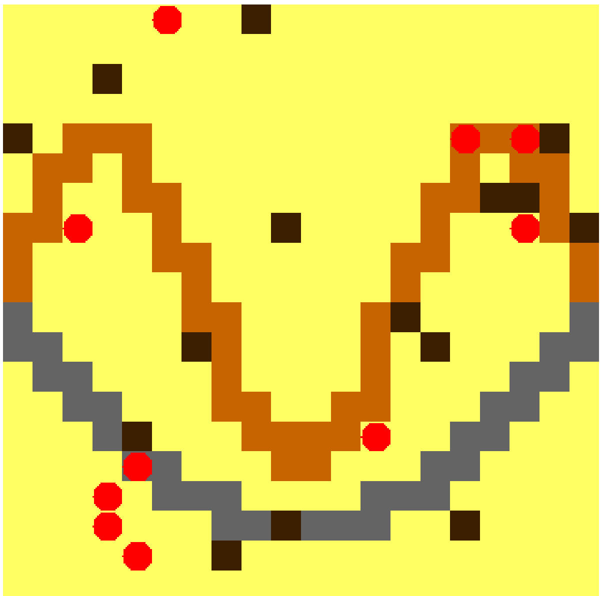

Figure 2: Our simple agent-based simulation. Grid cells are coloured yellow (sand), grey (the curving road), light brown (the meandering path), or dark brown (rocky areas). Red circles show IEDs, which are invisible to the Blue agents. The Blue agents begin at the middle left, and must travel to the right-hand side of the region, while avoiding the IEDs.

3. OPTIMISATION

Although the simulation shown in Figure 2 is very simple, it illustrates the basic dilemma faced by Blue forces in a Counter-IED context: which Blue strategy gives the best chance of survival? For experimental purposes, four Blue strategies were considered: a **direct** strategy of moving only to the right, and three terrain-based strategies (**sand**, **path**, and **road**) where Blue agents prefer to move along a specific kind of terrain, leaving it only when they are forced to.

In real life, Blue strategies include the selection of IED-resistant vehicles (JIEDDO, 2008; Zorpette, 2008a). They also include countermeasures to prevent the operation of infrared (Zorpette, 2008a, p 27) or radio triggers. An example radio countermeasure is the US CREW (Counter Radio-controlled-IED EW) system (JIEDDO, 2008, p 11).

Culvert denial systems such as the US Terrapin (JIEDDO, 2008, p 10) can also form part of a Blue strategy, as can detection devices, such as the US Fido (Zorpette, 2008a, p 29), or ground penetrating radar systems such as the US Husky (JIEDDO, 2008, p 10).

Less concrete Blue strategies include route planning; Standard Operating Procedures (SOPs), such as IED disposal techniques; intelligence-gathering activities for locating IEDS; and the full range of counterinsurgency (COIN) operations (US Army, 2006) which help to prevent IEDs from being placed.

Although our simple agent-based simulation does not capture the complexity of the full range of real-world strategies, it is sufficient to illustrate the basic optimisation approach to Operations Research. A range of Blue strategies is simulated, and results are plotted as in Figure 3.

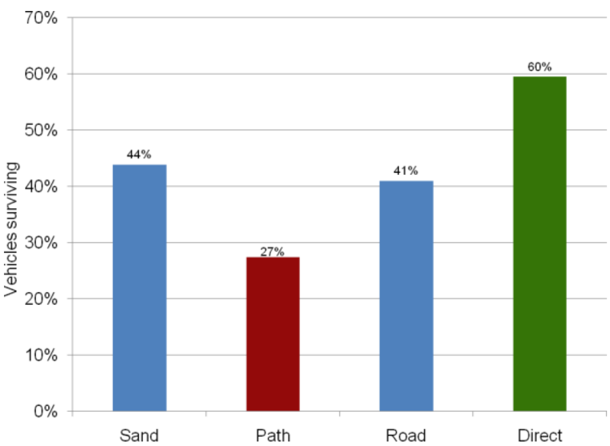

Figure 3: Performance of 4 Blue strategies against randomly-placed IEDs (averaged over 100,000 simulated Blue trips, each on a grid like Figure 2, but with different randomly-placed IEDs and rocky areas).

Where the space of Blue strategies is conceived of as two-dimensional, such a plot is called a "fitness landscape" (Ilachinski, 2004, p 503). If the range of possible Blue strategies is large, finding the optimum point on the fitness landscape may require sophisticated search techniques (Goldberg, 1989; Hecht-Nielsen, 1990).

In the case of Figure 3, however, it is clear that travelling along the meandering path is the worst strategy, with only 27.4% of Blue vehicles getting through, and the direct route is the best strategy, with 59.5% of vehicles getting through. These results are explained by the fact that, for this simple simulation, the shortest route is least likely to encounter an IED.

3.1 Weaknesses of the Optimisation Approach

The main weakness of this naive optimisation approach is that it assumes a single fixed Red (insurgent) strategy. IEDs are, by definition, improvised and ever-changing, at least over the longer term. Insurgents placing IEDs deliberately choose the strategies which they believe will be most destructive. The simplistic answers suggested by plots like Figure 3 are therefore inadequate, because they ignore the mind of the enemy.

4. GAME THEORY

Incorporating the enemy's ability to choose brings the problem into the domain of *game theory*. For experimental purposes, we provide six possible Red strategies: four terrain-based strategies (where IEDs are placed preferentially on **sand**, **path**, **road**, or **rocky** squares), one strategy where IEDs are placed preferentially in the **central** region of the grid, and one strategy where IEDs are placed **randomly** as they were in Section 2.

In real life, Red strategies include the type of IED (buried, explosively formed penetrator, etc.), the triggering device (radio, mobile phone, command wire, pressure plate, infrared sensor, etc.), placement options, camouflage options, and possible decoy devices.

In our simulated region, as in real life, the outcome of a Blue trip is a combination of the chosen Blue strategy for driving and the chosen Red strategy for IED placement. Table 1 shows the percentage of Blue vehicles getting through, averaged over 100,000 simulated trips. The row corresponding to Red's "random" strategy contains precisely the numbers in Figure 3.

Some of the combinations in Table 1 are particularly good for Blue. If Blue takes the road, for example, and Red places IEDs on the path, on the sand, or in the central region, then at least 92.5% of the vehicles will get through, with the exceptions resulting from cases where Blue strays off the road to avoid a rocky area. Conversely, if Blue and Red make the same choice of the path or the road, then less than 1% of the vehicles will get through.

Table 1: Percentage of vehicles getting through for combinations of Red and Blue strategies (averaged over 100,000 simulated Blue trips).

Red Strategies	Blue Strategies: Sand	Path	Road	Direct
Sand	37%	87%	93%	63%
Path	70%	0%	94%	41%
Road	66%	65%	1%	49%
Rock	81%	54%	66%	59%
Central	26%	10%	94%	19%
Random	44%	27%	41%	60%

Game theory assumes that Red and Blue choose their strategies independently, and takes account of the intelligence of both sides. Linear programming can be used to find a strategy (or probabilistic choice of strategies) yielding the best possible result against a totally rational opponent (Taha, 1992). Game theory was used in World War II for planning anti-submarine operations (Leonard, 1992), and in recent times has been applied to the IED problem as well (Washburn, 2006).

Applying the standard linear programming techniques to Table 1 yields the solution in Table 2. The optimal solution for both sides is a probabilistic choice of strategies (or "mixed strategy"). The Blue strategy guarantees an expected chance of getting through of at least 45.0%. The Red strategy guarantees an expected chance of getting through of at most 45.0%.

Table 2: Optimal solution for the game in Table 1, with each side making a probabilistic choice amongst three strategies. An expected 45.0% chance of vehicles getting through is the best that both sides can hope for against a totally rational opponent.

Red	Blue Strategies and probabilities: **Sand** 0.58	**Road** 0.29	**Direct** 0.12
Road 0.28	66%	1%	49%
Central 0.29	26%	94%	19%
Random 0.44	44%	41%	60%

4.1 Weaknesses of the Game Theory Approach

One weakness of the kind of analysis presented in Table 1 is that it assumes that the full range of possible Red and Blue strategies is known *a priori*, and that the percentages in the table are also known to both sides. In fact, being improvised, Red IED strategies are constantly being developed, as are Blue countermeasures against them and Red counter-countermeasures. Furthermore, the real-world equivalent of the percentages in Table 1 may be estimated ahead of time, but can only be known with certainty by experience. A solution like the one in Table 2 is therefore difficult to achieve in practice.

In addition, although the "mixed strategy" in Table 2 is expressed in probabilistic terms, it in fact applies only to a single trip, where neither side has any idea of the opponent's strategy. However, the counter-IED problem is an *iterated game*, involving a sequence of multiple trips. Information about the opponent's past strategies provides significant information about their future actions (even if the information is not totally certain). This is particularly true because neither side can switch strategies instantaneously. It takes time to develop and disseminate new Standard Operating Procedures (SOPs), and it also takes time to develop new equipment, such as new vehicles. We will examine the issue of iterated strategies further in Section 5.

4.2 Counterinsurgency and Nonzero-Sum Games

The analysis in Table 1 presents the counter-IED problem as a *zero-sum game*, where Blue is trying to minimise losses and Red to maximise them. While this may be adequate in the short term, in the long term counter-IED operations are part of the larger arena of counterinsurgency (COIN) operations, which are decidedly *nonzero-sum*.

The goal of counterinsurgency operations is a win/win solution where the insurgency diminishes because the concerns of potential insurgents are met. For example, during the Philippine communist insurgency of the 1950s, Defence Secretary (and later President) Ramon Magsaysay shrank support for the communist guerrillas by instructing blue forces to provide aid, medical assistance, and legal advice to villagers (Joes, 2008). Conversely, as the US Army/Marine Corps counterinsurgency field manual indicates, an aggressively zero-sum approach may be ineffective (US Army, 2006: 1-45), in that individuals who have been negatively impacted by counterinsurgency operations may join the insurgency, rather than support government forces (Chiarelli and Michaelis, 2005, p 6).

For nonzero-sum games, the equivalent of the optimal "mixed strategy" in Table 2 is a *Nash equilibrium*. However, as is well-known, a Nash equilibrium may not be *Pareto optimal*, that is, it may miss win/win solutions which are better for both sides (Poundstone, 1992; Morris, 1994, p. 127). Within nonzero-sum games such as counterinsurgency, mathematical analysis is often less important than developing an understanding of the social concerns of both sides, and creating an atmosphere of mutual trust in which a win/win solution can be accepted by a large majority of the participants. This may involve Operations Research focussed on network analysis and course-of-action ontologies (Darr *et al*., 2010).

5. ADAPTATION

An important aspect of IEDs is that they are, by definition, improvised and (in the long term) constantly changing. In response to Blue countermeasures, Red counter-countermeasures will be developed.

Even very sophisticated countermeasures may be negated by inexpensive or improvised counter-countermeasures (Zorpette, 2008a, p 30). Insurgents also adapt the design and placement of IEDs in response to Blue tactics (Zorpette, 2008b, p 38). Consequently, as the US Army/Marine Corps counterinsurgency field manual notes:

> "*Competent insurgents are adaptive. … Insurgents quickly adjust to successful COIN practices and rapidly disseminate information throughout the insurgency. … Effective leaders at all levels avoid complacency and are at least as adaptive as their enemies. … Constantly developing new practices is essential*." (US Army, 2006: 1-155)

It is difficult to explore innovation and adaptation in detail within a simple simulation model, but we have constructed a variation of the game-theoretic approach presented in Section 4, suitable for preliminary investigation. In this model, Blue chooses from the four strategies in Table 1, with a bias towards strategies that have worked well in the last 20 trips. Red makes a similar choice from the six Red strategies in Table 1.

Figure 4 shows the results of this simple form of adaptivity. The number of Blue vehicles getting through oscillates between 28% and 70%, with a mean of 48.8%. Auto-correlation analysis suggests that Figure 4 consists of random noise overlaid on an irregular oscillation with a period of about 600 trips. In other words, after about 300 trips Red (or Blue) is able to detect and exploit a pattern in the other party's actions, and temporarily gain a slight upper hand. After another 300 or so trips, Blue (or Red) will in turn find a successful counter-response. The alternation between Red and Blue advantage continues throughout this iterated game.

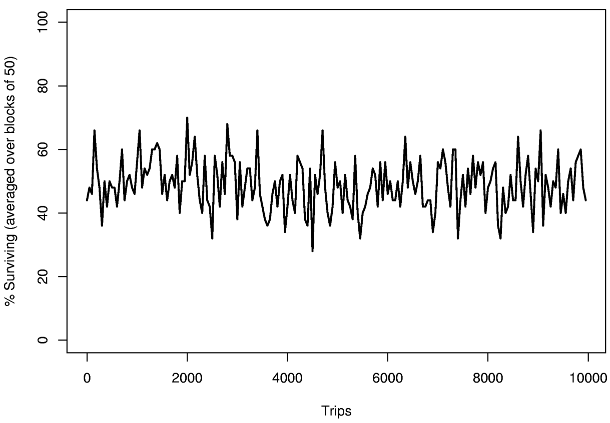

Figure 4: Results of Red and Blue adaptation over 10,000 trips (averages over groups of 50 trips).

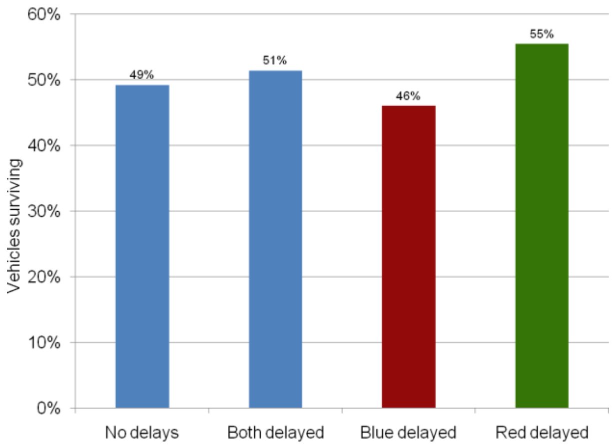

Figure 5: Results for reducing Red and/or Blue adaptivity by a delay factor. The delayed party does significantly worse (differences between colours statistically significant at the 10^{-15} level or better).

Figure 4 reflects a situation where Red and Blue are equally adaptive. We varied this in a further experiment by giving Red and/or Blue a 100-trip delay in responding to past events. In other words, instead of using strategies that worked well for trips at time t–20 to t–1, the delayed party responded to trips at time t–120 to t–101. Figure 5 shows the results. When both sides are equally adaptive, around 50% of vehicles get through, but this increases to 55.5% when Blue is more adaptive, i.e. Blue is "inside the OODA[†] loop" of Red (Brehmer, 2005). It drops to 46.1% when Red is more adaptive.

6. EVOLVING NEURAL NETWORKS

A somewhat more realistic adaptivity model allows Red and Blue to innovate new strategies using *genetic algorithms* (Goldberg, 1989) to evolve *neural networks* (Beale & Jackson, 1990; Hecht-Nielsen, 1990; Nowak *et al*., 1998). This evolutionary process is similar to that of Floreano *et al*. (2001) or Klüver & Klüver (2010). It provides a better model of human learning processes than that given above, and one which is applicable to more than just the "toy" problem of Figure 2. We have constructed a simple Java-based evolutionary system where Red and Blue can develop decision-making approaches which generalise the strategies of Table 1.

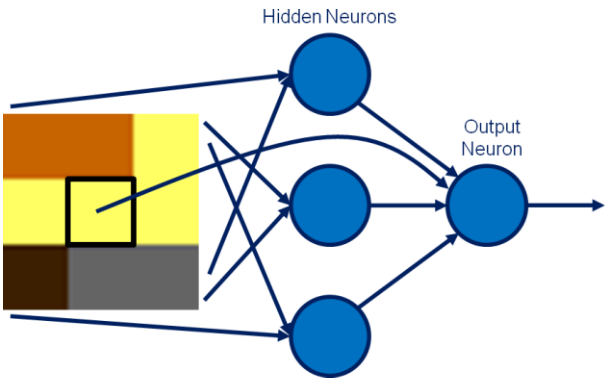

Figure 6: Decision-making using a neural network. Nine terrain squares feed into the hidden neurons. The central square also feeds directly to the output neuron.

[†] Observe-Orient-Decide-Act

First, Red and Blue use neural networks for decision-making, as in Figure 6. For Red, the input to the neural network is an encoding of the terrain at and around a candidate IED placement point. For Blue, the input is an encoding of the terrain at and around the current position, together with an encoding of past actions. Neurons apply a sigmoid function to the sum of their weighted inputs, and the result is discretised to obtain as a binary placement decision (for Red) or a left-ahead-right movement decision (for Blue).

Each possible neural network is defined by a list of its weights (31 weights for Red and 40 for Blue), which in turn can be specified by a bit string. Genetic algorithms are used to evolve bit strings specifying effective neural networks. As the opposition changes, the definition of "effective" changes also, and so the populations of Red and Blue neural networks co-evolve.

Figure 7 shows the result of the evolutionary model. Blue vehicles getting through oscillate between 28% and 68%, with a mean of 46.3%. Oscillations have a period similar to that of Figure 4.

The similar behaviour indicates that the simpler adaptivity model was in fact adequate for drawing the conclusions that were made: that adaptivity leads to oscillations, as first one side and then the other gains an edge; and that the side which adapts more rapidly has an overall advantage. For more complex problems, we would expect the benefits of the greater realism of evolved neural networks to become apparent, in spite of their inevitably greater demand on computer resources.

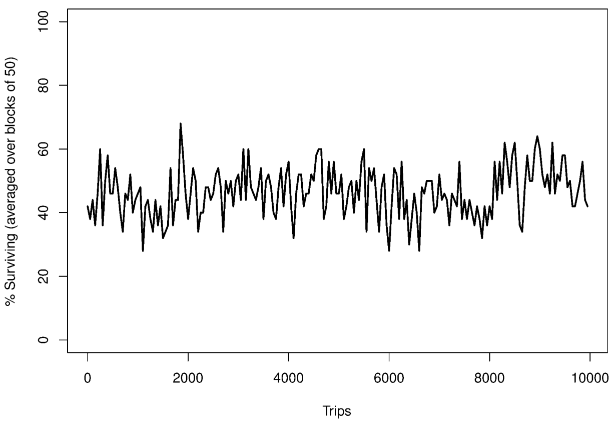

Figure 7: Results of Red and Blue adaptation over 10,000 trips using evolution of neural networks (averages over groups of 50 trips).

Figure 8 shows the result of making one or other side less adaptive, by adding noise to the performance measures used by that side during evolution. As in Figure 5, the less adaptive side succeeds less often, overall.

Investigating the extent of this advantage in the real world would, of course, require a more realistic model than the one in Figure 2. However, our results suggest that genetic algorithms are a promising way of studying real-world adaptivity in such a more realistic model.

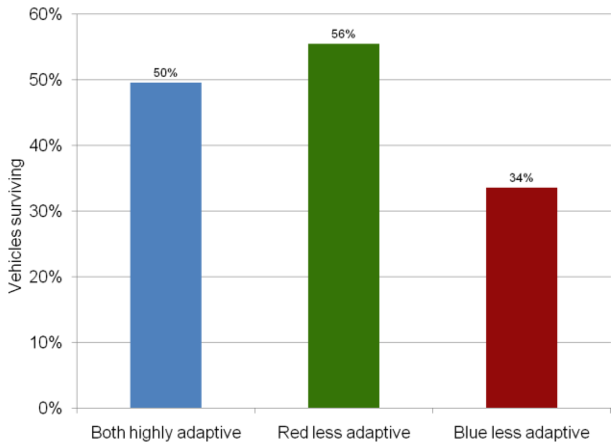

Figure 8: Results for reducing Red or Blue evolutionary adaptivity (averages over 20 runs each). The less adaptive party does worse (differences statistically significant by ANOVA at the 10^{-15} level or better).

7. DISCUSSION

In this paper, we have used the simple agent-based simulation in Figure 2 to briefly survey Operations Research issues associated with counter-IED activities.

While it is possible to assume a fixed Red strategy and optimise against it, as in Section 3, this fails to capture the improvised nature of IEDs. While the game-theoretic approach in Section 4 captures the fact that Red is an intelligent opponent, it still does not recognize the ever-changing nature of the IED threat, and the nonzero-sum aspect of counterinsurgency operations.

Doing justice to the IED threat requires incorporating adaptivity into the model, so that Red or Blue are, in a sense, optimising on a fitness landscape which constantly changes as the opponent adapts. Figure 5 and Figure 8 highlight the fact that in such a contest, the most rapidly adapting side has an advantage. This conclusion reinforces other work on adaptive learning (Spaans *et al.*, 2009). Our results suggest that genetic algorithms for evolving neural networks are one promising way of studying real-world adaptivity.

REFERENCES

Beale, R. & Jackson, T. (1990) *Neural Computing: An Introduction*, Adam Hilger.

Brehmer, B. (2005) "The Dynamic OODA Loop: Amalgamating Boyd's OODA Loop and the Cybernetic Approach to Command and Control," *10th International Command and Control Research and Technology Symposium (ICCRTS)*: www.dodccrp.org/events/10th_ICCRTS/CD/papers/365.pdf

Chiarelli, P.W. & Michaelis, P.R. (2005) "Winning the Peace: The Requirement for Full-Spectrum Operations," *Military Review*, July-August, pp 4–17. www.au.af.mil/au/awc/awcgate/milreview/chiarelli.pdf

Darr, T., Benjamin P., & Mayer, R. (2010) "Course of Action Ontology for Counterinsurgency Operations," *15th International Command and Control Research and Technology Symposium (ICCRTS)*: www.dodccrp.org/events/15th_iccrts_2010/papers/126.pdf

Davis, M. D. (1983) *Game Theory: A Nontechnical Introduction*, Basic Books.

Floreano, D., Nolfi , S., & Mondada , F. (2001) "Co-Evolution and Ontogenetic Change in Competing Robots," in Patel, M., Honavar, V., & Balakrishnan, K. (eds), *Advances in the Evolutionary Synthesis of Intelligent Agents*, MIT Press.

Goldberg, D.E. (1989) *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley.

Hecht-Nielsen, R. (1990) *Neurocomputing*, Addison-Wesley.

Ilachinski, A. (2004) *Artificial War: Multiagent-Based Simulation of Combat*, World Scientific.

JIEDDO: Joint Improvised Explosive Device Defeat Organization (2008) *Annual Report*. https://www.jieddo.dod.mil/ANNUALREPORTS/20080130_FULL_Annual%20Report%20UNCLS%208.5%20x%2011_v6.pdf

Joes, A.J. (2008) "Counterinsurgency in the Philippines 1898–1954," in Marston, D. & Malkasian, C. eds., *Counterinsurgency in Modern Warfare*, Osprey.

Klüver, J. & Klüver, C. (2010) *Social Understanding: On Hermeneutics, Geometrical Models and Artificial Intelligence*, Springer, Chapter 3.

Leonard, R.J. (1992) "Creating a Context for Game Theory" in *Toward a History of Game Theory*, E.R. Weintraub, ed., Duke University Press, pp. 64–65.

Morris, P. (1994) *Introduction to Game Theory*, Springer.

Nowak, A., Vallacher, R. R., & Burnstein, E. (1998) "Computational social psychology: a neural network approach to interpersonal dynamics" in Liebrand, W., Andrzej Nowak, A., & Rainer Hegselmann, R., *Computer Modeling of Social Processes*, Sage, pp. 97–125.

Poundstone, W. (1992) *Prisoner's Dilemma*, Oxford University Press.

Spaans, M., Spoelstra, M., Douze, E., Pieneman, R., & Grisogono, A.-M. (2009) "Learning to be Adaptive," *14th International Command and Control Research and Technology Symposium (ICCRTS)*: www.dodccrp.org/events/14th_iccrts_2009/papers/005.pdf

Taha, H.A. (1992) *Operations Research: An Introduction*, 5th ed, Macmillan, chapter 12.

US Army (2006) *Field Manual 3-24 (US Marine Corps MCWP 3-33.5): Counterinsurgency*. usacac.army.mil/cac/repository/materials/coin-fm3-24.pdf

Washburn, A. (2006) *Continuous Network Interdiction*, Report NPS-OR-06-007, Naval Postgraduate School, Monterey, California.

Zorpette, G. (2008a) "Countering IEDs," *IEEE Spectrum*, Vol. 45, No. 9, September, pp 22–31.

Zorpette, G. (2008b) "Bomb Squad Diary," *IEEE Spectrum*, Vol. 45, No. 10, October, pp 34–41.