SUPPORTING THE ARMY CAPABILITY DEVELOPMENT PROCESS USING AGENT BASED DISTILLATIONS – A CASE STUDY

DR ANDREW GILL
MSEB, DSTO

DR RICHARD EGUDO
LOD, DSTO

DR PETER DORTMANS
LOD, DSTO

MR DION GRIEGER
LOD, DSTO

Prepared for submission to the Journal of Battlefield Technology
(www.adfa.edu.au/jbt)

SUPPORTING THE ARMY CAPABILITY DEVELOPMENT PROCESS USING AGENT BASED DISTILLATIONS – A CASE STUDY

Andrew W Gill, Richard Egudo, Peter J Dortmans and Dion Grieger

Abstract. The Australian Army is seeking information addressing the combat effectiveness of their conceptual Enhanced Combat Force (in a 2015 timeframe). One of the main questions to be answered is how their manoeuvre concepts might need to change. To experiment with agent based distillations we abstracted a problem based on Manoeuvre Operations in a Littoral Environment and the specific hypothesis to be tested was whether a small, mobile force with high situational awareness coupled with effective reach-back munitions could defeat a significantly larger force. The aim of this paper is to illustrate the application of EINSTein and present some preliminary results against our hypothesis, and to make some initial observations on the potential utility of such models for land operations analyses.

INTRODUCTION

Background

In order to develop a "dynamic and evolutionary" war-fighting capability and in response to the Revolution in Military Affairs, the Australian Army initiated a process for remodeling the Army that would both enhance its current capacity to meet its strategic requirements, and provide direction for the migration to an enhanced future combat force. Known as 'Restructuring the Army' (RTA), it commenced field trials in 1997, in order to analyse, develop and enhance capabilities and processes, and provide evidence to inform decisions on the types of capabilities Australia should invest in, in the medium to long term [1].

A central component of this methodology was the Battlelab Process, which focused on modeling systems, testing them in the field and then analysing those results with the view to further refine that model. This process was further refined and embedded within the Army Experimental Framework (AEF) [2], which provided a 6 step process for military experimentation. The RTA trials were underpinned by a vision based on the manoeuvre concept, that is an integrated modern highly mobile task forces and units capable of effective autonomous operations of widely dispersed and dynamic nature in both joint and combined theatres.

The most difficult, resource intensive and time-consuming phase of RTA were the field trials conducted in 1998. Therefore, for Phase 2 (the Headline Experiment 99 (HE99)) the experiment utilised seminars and war-games rather than field trials and focused on determining the impact of varying levels of situational awareness on an austere, highly mobile but organically firepower poor force fighting in open terrain. The HE99 experiment itself involved considerable effort from both the defence and scientific communities in the design, conduct and analysis of the 2 week experiment.

The results were later fed into higher resolution war-games and closed loop simulations. Coding the scenarios took 3 months, so the preliminary analysis results became available six months after the experiment was completed. However, as AEF activities are an annual event, planning for HE00 was already underway, so that some opportunities for further refinement of the concepts were missed.

This highlights the high resource and time requirements that current land combat analysis tools require in providing results to inform capability development decisions. Lauren and Baigent [3] also outlined other difficulties traditional war-games and simulations have with modelling land-force issues, which has led them to investigate alternative models under the Project Albert research program.

Project Albert is a United States Marine Corps (USMC) research effort that attempts to assess the general applicability of the concept of 'Operational Synthesis' [4] to land warfare. Project Albert aims to identify emergent behaviour through the application of a bottom-up rather than top-down approach and seeks to address three key areas: non-linear behaviour (where small changes create disproportionate responses); co-evolving landscapes (which characterise the changing battlefield) and intangibles (such as morale, discipline and training) for which conventional land combat analysis models are particularly poor at investigating.

The NZ Defence Technology Agency has been active recently in using the tools within Project Albert to assist in restructuring their combat force. The Australian Army and DSTO have subsequently become collaborators within the Project Albert research program.

Agent Based Distillations

Agent based distillations (ABD) are low-resolution abstract models, used to explore questions associated with land combat operations in a short period of time. Being agent based means that only simple behavioural rules need to be assigned. This is generally achieved by assigning 'personalities' to the agents by way of relative weightings to various elements on the battlefield (friendly and enemy agents, notional 'flags', terrain features, etc) and a linear penalty function to determine the entity's next move. Various 'meta-personalities' can also be assigned which moderate the agent's default personality if certain threshold constraints are exceeded from time to time. Being deliberately low-resolution means that the detailed physics of combat are largely ignored (or abstracted to simple constructs).

Thus the scenario is much less scripted than that of traditional war-games, the idea being to allow a focusing of thought on the essential elements of the systems, which typically is the dynamic interaction of entities on the battlefield. Advances in computing power can then be exploited to produce a significant volume of data. This process is known as data farming [5] and allows extensive parameter excursions to be performed, both in terms of variations in platform capabilities and tactics (behavioural characteristics), from the baseline scenario. This then enables multi-way sensitivity analyses to be performed to explore any non-linear behaviour and synergies in the system. The farmed data can also be used to perform statistical analyses to test the significance of the properties observed.

This is in stark contrast with traditional war-games whose timescales are measured typically in units of weeks or months. The trade-off is that the modeling resolution using ABD is sacrificed. Thus the level of abstraction implies that the results of a distillation should only be used to provide a focusing of ideas and that subsequent analyses be conducted to 'drill-down' with higher resolution modeling.

There are a growing number of ABD being used under Project Albert, including the Irreducible Semi-Autonomous Adaptive Combat (ISAAC) model [6] and the Enhanced ISAAC Neural Simulation Toolkit (EINSTein) [6]. The NZ DTA has also recently developed the Map Aware Non-uniform Automata (MANA), to support their studies [7].

The Case Study

HE99 was designed to provide information addressing the combat effectiveness of a 2015 Enhanced Combat Force. A central question was how Army's manoeuvre concepts might need to change. To experiment with ABD we abstracted a problem based on Manoeuvre Operations in a Littoral Environment (MOLE) and the specific hypothesis to be tested was whether a small, mobile force with high situational awareness coupled with effective reach-back munitions could defeat a significantly larger force.

A 3-day workshop investigated this proposition employing the EINSTein [6] distillation to facilitate the study. The workshop had three aims. First, a number of baseline scenarios were to be constructed which modelled the units and mission as best could be achieved. As a result of this process, two subsequent aims should also have been achieved. They are, to determine some of the limits of applicability and resolution of the EINSTein distillation in modelling or representing Army capabilities and missions, and to develop within the CATDC-DSTO group an increased level of proficiency in the use of ABD.

MOLE SCENARIO

The main physical characteristics of each element are presented in Table 1. The Blue force consists of a mix of light armoured vehicles (LAV), armed reconnaissance helicopters (ARH) and HIMARS. For the baseline scenario, the force mix is such that there are 10 LAV, 5 ARH and 1 HIMARS unit, while the Red force consists entirely of tanks (45 T-80's). Thus, the Red to Blue force ratio is approximately 3:1.

	LAV	ARH	HIMARS	T-80
Speed	2	4		1
Sensor	4	8		2
Fire	2	4		2
Lethality	0.25	0.5	0.75	0.5
Number	10	5	1	45

Table 1. Major Physical Characteristics

The LAV have relatively good speed and sensor range, but relatively poorer weapon characteristics. The task for the LAV is to survey the likely approaches of the enemy and to communicate detections back to the ARH and HIMARS units for prosecution. The ARH are significantly faster than the LAV and have double their sensor and weapon performance, however there are fewer of these assets. The task for the ARH is to quickly move to the location of detected enemy and decisively engage, based on the communicated information supplied by the LAV. The HIMARS unit is a single asset held at the rear of operations and brings heavy, lethal area-fire onto regions of detected enemy supplied by the LAV. The T-80 has half the movement and sensor characteristics of the opposing LAV, but have double the weapon performance and out number the LAV.

To simulate reconnaissance behaviour, 'negative attractiveness' to friendly and enemy entities is used. The former is used to create a dispersed reconnaissance force, while the latter is used to ensure the LAV does not become decisively engaged. A high attractiveness to the Area entity is used to simulate an area of operations (AO) assigned to the LAV force. The Cluster 'meta-personality' was also used to further enhance the dispersed nature of the LAV force. Similar entity definitions can then be constructed for the other units (ARH, HIMARS, T-80) to simulate the required characteristics and behaviours.

HIMARS Modeling

HIMARS proved the most difficult entity to represent. EINSTein does not explicitly model indirect or area fire weapons, in particular the forward observer concept. The closest approximation was to assign a grenade to a HIMARS squad that was given a low sensor range and a high movement range to allow it to quickly react to communicated information. That is, the HIMARS would actually move quickly to where the target is and when it was within its limited throwing range it would fire a munition. When no enemy agents were present in it's sensor range and no information was being received from the forward observers it would then quickly retreat to its initial position.

The problem using this representation is that enemy agents would react to the HIMARS when it was in their sensor range. Of course ideally the HIMARS would be located stationary at the rear but it was hoped that the high movement range and ability of the HIMARS to advance and retreat so quickly would minimise this unwanted behaviour.

The grenade weapon parameters used for the HIMARS are shown in Figure 1. You will notice that the Probability of Hit may seem relatively low (0.4 as compared to 0.5 for the Red tanks). Another feature that cannot be modelled directly in EINSTein is a time lag between rounds fired. A weapon such as HIMARS requires a non-insignificant time between rounds to reload and acquire a target. It was found that a high Probability of Hit value for HIMARS was too lethal, and that the lower value of 0.4 provided more realistic behaviour and could be viewed as a form of time delay between rounds.

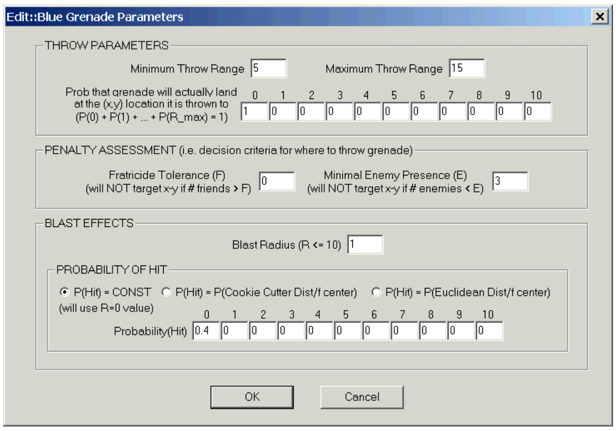

Figure 1. Grenade Parameters for HIMARS

RESULTS

Interactive Playback Mode

This enables the analyst to examine the behaviour of the entities, which should be correlated with their desired characteristics and tasks. A degree of fine-tuning of the entity parameters is generally required to produce a baseline scenario with all entities functioning in a representative and consistent way. However, one should try to avoid tweaking the parameters unnecessarily in an effort to produce the 'correct behaviour', that is, to produce scripted behaviour. The central point of ABD's is to seek emergent behaviour from the local interaction rules we define – not to constrain that behaviour.

Once the fine-tuning has been performed and a baseline scenario constructed, the Interactive Playback mode allows the analyst to obtain qualitative information about the force mix dynamic interactions. For our baseline scenario, Red travels tightly grouped from East to West through the AO patrolled by the LAV squad. The LAV, due to their superior sensors and speed, detect the incoming T-80 and communicate these detections back to the waiting ARH and HIMARS. From the ensuing engagements we note that most LAV manage to avoid decisive engagement with the T-80 and generally survive. The Red force is heavily attrited, mainly by the ARH and HIMARS and only a few Red manage to reach the objective (represented by the Blue flag).

Thus for the baseline scenario, at least on a qualitative level, it is not impossible for a smaller, more mobile force with high SA and effective reach-back munitions to defeat a much larger opposing force. The question that arises is what is the relative contribution to this success of differing force mixes and varying asset characteristics.

One Way Sensitivity Analysis

This allows the relative effect of individual parameters on the mission to be quantified. As an example of this parameter excursion, we investigated the effect of different force mixes (in terms of the number of ARH and whether or not HIMARS was available) on the success rate of the Red force. The measure of effectiveness (MOE) used was the percentage of Red forces that manage to reach the objective (Blue flag).

Figure 2 shows the variation of this MOE with different numbers of ARH – the upper curve represents the situation with no HIMARS while the lower curve is the case with a single HIMARS unit. With no HIMARS and no ARH the Red force easily achieves its mission, with all entities reaching the objective. With a single HIMARS and no ARH just over half of the Red force now manage to reach the objective. In both cases, as the number of ARH is increased, Red mission success is diminished.

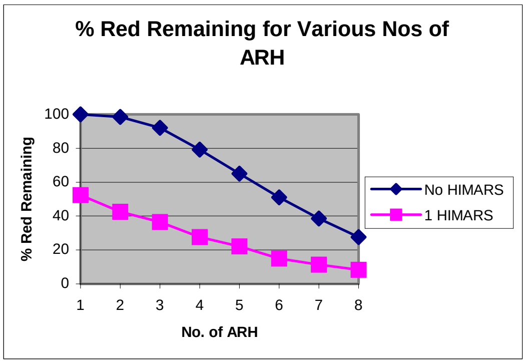

Figure 2. Snapshots of Baseline Scenario Simulation

In both cases, there is some non-linearity present in this diminishment, though it is not strong. In the case of no HIMARS, it appears that at least two ARH are required to significantly affect Red's mission. Also, in the case with HIMARS, there appears to be diminishing returns as more and more ARH are added to the force mix. This may suggest that there is an upper limit of ARH that a cost-effective Blue force mix should possess.

This graph can be used to make capability comparisons. For example, the data indicates that to ensure that only 50% of the Red force achieves their objective, this effect could be equally generated with either one HIMARS or six ARH. Similarly, to ensure that only 30% of the Red force achieves their objective, this effect could be equally generated with either one HIMARS with four ARH or eight ARH. Note that this second result does not scale linearly with the first (which would suggest that one HIMARS with four ARH is equivalent to ten ARH). This type of force mix trade-off analysis could be useful in supporting acquisition decisions once the relative costs of assets are taken into account.

Fitness Landscape

Essentially, this is a 2D sensitivity analysis and the surface plotted shows the variation of the selected MOE with two user-specified parameters, which is a useful mechanism to detect allowable trade-offs (essentially contour lines of the plotted surface) as well as synergies between parameters.

Figure 3 examines the variation of the Red to Blue Survival Ratio (a complement to the usual loss exchange ratio (LER)) to changes in the size of the Red force (ranging from 15 to 60) and to the level of dispersion of the Red entities (ranging from low to high). The latter was modelled by using the Minimum Distance to Friendly meta-personality. Higher values of the MOE indicate improved Red mission success.

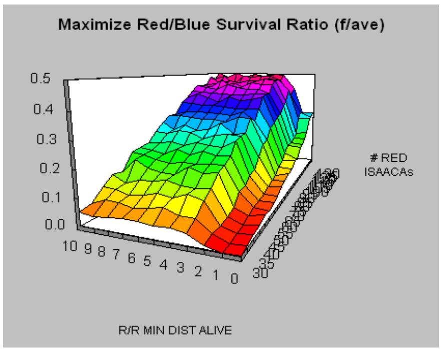

Figure 3. Force Size and Dispersion Level

If we take slices of the surface for different dispersion levels, the shape of the curve is roughly linear with the number of Red forces. Thus, combat weight for Red appears to have a linear effect on success. The surface also clearly shows marked improvement for Red once a dispersion level greater than 1 is achieved. For dispersion levels greater than 3, for a fixed force size, there in no noticeable improvement. Thus, the optimum dispersion level appears to be roughly 3.

The cause of this result was deduced by running several Interactive Playback sessions, which reveal that the reason is related to the means of employment of the HIMARS. As HIMARS is a limited resource, thresholds were imposed such that delivery of a HIMARS round required a minimum number of enemy targets within a given range and a maximum number of friendly entities. Thus, once Red dispersed to a certain level, it effectively provided Blue with no sufficiently massed target to afford a HIMARS strike by remaining below its engagement threshold.

This result immediately suggests 'what-if' scenarios and ABD's can be used to game these combinations. As mentioned above, this Fitness Landscape analysis can allow trade-offs to be explored. For example, it might be possible for Red to use a smaller but more dispersed force and achieve the same level of mission success.

Figure 4 below displays the Fitness Landscape when varying the sensor range and probability of kill (lethality) of Red. Once again, if we examine slices of this landscape for fixed values of the sensor range, we see that the lethality of Red appears to have a linear effect on its mission success.

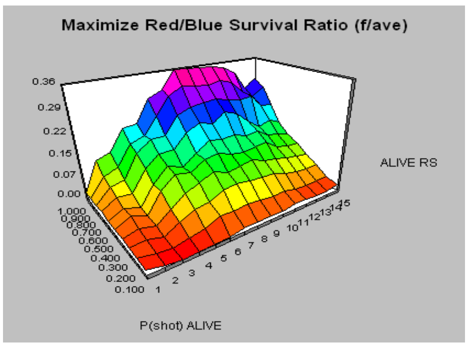

Figure 4. Sensor Range and Probability of Kill.

However, the interesting point to note is that the degree of linear effect (essentially the slope of the curve) is not constant but changes quite strongly as the sensor range of Red is increased. Initially this change is positive, whereby the effect of an increase in lethality from 0.4 to 0.6 (for example) is more pronounced with a sensor range of 6 than with a sensor range of 2. This illustrates the potential effect of synergy between platform characteristics.

Note also, however, that this behaviour does not occur for all values of the sensor range, and in fact a reversal of behaviour appears to occur once a sensor range of about 8 is exceeded. On further investigation (by using the Interactive Playback Mode) the cause for this behaviour was traced to the termination criteria of the simulation that produced unrealistic behaviour in those cases.

The goal for the Red force is to reach the Blue objective (the flag) while attempting to minimise it's own losses and maximizing losses to the Blue force. The termination criteria used to stop the simulations and collect data on force losses was reaching a fixed time, which needs to be set large enough to allow the mission to be played out. In most cases, the Red force made its way to the objective where it then waited safely until the termination time was reached. However, in the cases where it's sensor range was large, it could detect the Blue forces and was drawn back into battle and away from its objective, and suffered increased losses as a result.

Thus the results for these cases should be discarded. However, this analysis is useful in highlighting the need to critically examine the data output and it's relevance to the problem under investigation, and the Interactive Playback mode is a useful tool to achieve this.

Again one can also use these landscapes to trade-off parameters, whereby for example the same effectiveness for Red is achieved with a sensor range of 2 and a probability of kill of 1 or a sensor range of 5 and a probability of kill of 0.4. One might suspect that the technological challenges of achieving such a high lethality in the former configuration are such that the latter solution might be more feasible.

Dispersion versus Speed

A final trade-off analysis conducted for this scenario was that between the speed of the Red force tanks and the level of dispersion adopted. From the Interactive Playback runs, it is apparent that the casualties suffered by Red occur in the time taken to traverse from its staring position to the objective on the West side of the battlefield. If that time taken could be reduced, then Red would expect to take fewer losses on average.

Thus the situation considered was one of a choice for Red to either conduct its movement along a road or cross-country. The effect of road travel was to increase the speed of the tanks but at the expense of having to travel in a more grouped (or less dispersed) fashion. Cross-country travel was slower but could be performed at different levels of dispersion. Due to the limited number of movement speeds within EINSTein, the speed improvement of on-road travel was taken to be a doubling of the cross-country speed.

EINSTein was used to produce loss exchange ratio (LER) data under three situations – cross country with low dispersion; cross country with medium dispersion; and on road (therefore with no dispersion). Table 2 below displays the results generated. Note that a larger LER value corresponds to improved Red performance.

	Low Dispersion	Medium Dispersion	On Road
Red Losses	91%	66%	68%
Blue Losses	27%	50%	26%
LER	0.30	0.76	0.38

Table 2. LER for Different Modes of Red Movement

The results indicate that dispersed travel is preferable if travelling cross-country (which is essentially what the Fitness Landscape in Figure 2 above revealed), in that both Red losses are reduced and Blue casualties are increased and the LER is consequently more than doubled. The results also indicate that if travelling on road, then only the Red losses are reduced (by the same margin as dispersed cross country) but the Blue casualties are not affected. This is because of the decreased time Red has to engage the Blue LAV due to the increased speed on-road, and the decreased ability to hunt the Blue LAV due to being constrained to the road. Consequently there is only a marginal improvement in the LER.

Thus, if only the number of Red losses is important, then both tactics of cross-country dispersed or on-road travel are equally effective. However, if the LER is more important, then the results indicate that the tactic of cross-country dispersed travel would be preferable.

SUMMARY AND CONCLUSIONS

The case study analysed here produced a number of useful initial insights into the force mix problem. First, analysis by ABD's allowed quite quickly the contributions of the ARH and HIMARS assets to mission success to be quantified and traded off. The results suggested some regions of non-linearity (decreasing returns) for the ARH effectiveness and highlighted the importance of tactical considerations employed by the Red force against indirect weapons and allowed various tactical options to be evaluated including cross-country or route movement decisions.

Synergies among platform or weapon characteristics are easily identified using the Fitness Landscape run-mode, and for the force mix problem it was found that sensor range and lethality act quite strongly together. The implication is that investments in weapon and platform upgrades might be best considered jointly rather than in isolation.

Agent based distillations have potential for distilling a problem into the essential elements of the analysis, assuming these components can be modelled to the resolution required of the study. Parameter excursions can easily be conducted (either on PC's running overnight for more reliable statistics, or within say an hour for coarse grained results). This is in stark contrast with traditional war-games whose timescales are measured typically in units of weeks or months.

However, it was found that the EINSTein ABD did possess a number of undesirable characteristics. The code is somewhat unstable and some functionality that would have been very useful for the force mix hypothesis studied was either unavailable or did not function properly. For example, the modeling of indirect fires is very limited and some of the purported features associated with terrain did not function as described. Also, there were some variables that would have been quite useful if they were made squad specific, for example, communications range and the selection of targets and the associated lethality against that target.

Finally, it is important to stress that the results of a distillation merely provide some potential directions for further study, which may or may not prove to be useful (depending on the degree of abstraction required to 'fit' an ABD scenario). They do not provide quantitative 'answers'. Their usefulness, if proven to be true, lies in their ability to quickly provide a focusing of ideas for further higher resolution modeling (for example, in suggesting which factors appear to be important in subsequent war-gaming).

ACKNOWLEDGEMENT

The authors thank Colin Stanford, Peter Williams, Neville Curtis of LOD and MAJ Trevor Colton of CATDC.

REFERENCES

[1] D. Bowley, Restructuring of the Army Task Force Trials Scientific Framework Paper, DSTO Report, DSTO-GD-0169, 1999 (Restricted)
[2] The Army Experimental Framework, CATDC 1999.
[3] M. Lauren and D. Baigent, Exploring the value of sensors to a Recce Unit using Agent Based Models, Journal of Battlefield Technology, Vol 4, No 1, pp 1-6, 2001
[4] A. Brandstein, Operational Synthesis: applying Science to Military Science, Phalanx, Vol 32 No 4, 1999
[5] A. Brandstein and G. Horne, Data Farming: A meta-Technique for Research in the 21st Century, Maneuver Warfare Science 1998, Marine Corps Combat Development Command Publication.
[6] A. Ilachinski, Irreducible Semi-Autonomous Adaptive Combat (ISAAC): An Artificial-Life Approach to Land Combat, Military Operations Research, Vol 5, No. 3, pp 29 – 46, 2000.
[7] M.K. Lauren and R.T. Stephen, Using Fractals to Understand the Role of Entropy in Complexity Science: An Examination of the MANA Combat Model, Fractals (in preparation).

Dr Andrew Gill is a research scientist (RS) with the Military Systems Experimentation Branch of the Defence Science and Technology Organisation (DSTO), Australia. His research interests are in exploratory analysis techniques including agent based distillations and applied mathematics generally.

Richard Egudo obtain an Msc from London School of Economics and PhD from LaTrobe University. He was senior lecturer at Monash University prior to joining DSTO. He has published many articles in international journals mainly on mathematical programming.

Peter Dortmans obtained a PhD in Theoretical Physics at the University of Melbourne in 1992 and undertook post-doctoral research at the University of Melbourne and the Instituto Nazionale di Fisica Nucleare in Padova (Italy), before joining DSTO in 1999. He was subsequently posted as DSTO's Chief Analyst to the Australian Army.

Dion Grieger is an Industry Based Learning Student at LOD, DSTO. He is currently studying at the University of Adelaide and will be completing a double major in Applied Mathematics and Computer Science in 2002.