

# Suppression of Dismounted Soldiers: Towards Improving Dynamic Threat Assessment in Closed Loop Combat Simulations

**J. Millikan, M. Wong and D. Grieger**

*Land Operations Division, Defence Science and Technology Organisation  
Email: Martin.Wong@dsto.defence.gov.au*

**Abstract:** Closed loop combat simulations are part of a suite of analytical tools used to help inform decision makers about various military options and potential acquisitions. A common limitation of many current combat simulations is that entities often use scripted or time-based behaviours which limit their ability to imitate the dynamic nature of combat. In addition these behaviours cannot typically be used to simulate other military environments or scenarios without resource intensive input from both military subject matter advisers and simulation experts.

In this paper we attempt to address some of these limitations by applying and adapting techniques being used in both the computer gaming industry and the military research community. We propose a conceptual framework for modelling suppression in combat simulations that incorporates dynamic cognitive assessments of known and perceived battlefield threats, perceived cover locations, physiological state and social forces such as morale and leadership. We create a representative implementation of key aspects of the model in an agent-based test environment and apply it to a simple military scenario. We fuse information based on enemy lines of sight, incoming enemy fire, friendly lines of fire and the perception of dangerous areas of the environment to generate an individual “danger map” for each entity. An individual suppression model and social pressures based on simple team rules are incorporated with the danger map to generate dynamic entity behaviours.

We create a baseline behaviour type that is representative of that used in current closed loop simulations and progressively add dynamic behaviours based on the model described above. The results show that the more sophisticated dynamic behaviours increase force survivability while producing battles that are longer and provide a better representation of the ebb and flow of many combat situations.


**Keywords:** *Suppression, Defence, Military, Simulation, Operations analysis*

## 1. INTRODUCTION

Closed loop simulations of land warfare are often employed as part of an analytical campaign to gain statistically significant results on Force options from scenarios that are tightly constricted to minimal tactical or behavioural changes during each run. The campaign of analysis leading to such use often utilises computer-based human-in-the-loop wargaming where the behaviours of entities representing soldiers or civilians are controlled by military subject matter experts. The data captured during these wargames contributes to the generation of entity behaviour sets for closed loop simulations (eg CAEn (Shine and Coutts, 2006), Combat XXI). These behaviour sets currently do not autonomously appraise the spatial threat of the environment/enemy complex which can lead to undesirable entity behaviours occurring when the outcomes of stochastic events differ from the initial wargame and severely limits their ability to be used in other scenarios.

Current combat simulations also have a limited representation of the internal visceral reactions to enemy threats (eg due to survival instinct). A primary indicator of behavioural change due to these visceral factors is the suppression of the individual due to enemy fire or the threat of enemy presence. Suppression may be defined as "...the causing of human reactions that reduce individual (unit) efficiency to fire, observe and move."(Winter & Clovis, 1973). In the widest sense, suppression is the primary influence on how individuals act in the presence of potential or current threat. Perception of threat itself can be represented as a three-dimensional field of scalar threat values that each individual maintains as a result of their situational awareness and beliefs about how the environment advantages or disadvantages enemy action.

We address these limitations through the employment of an Agent-Based Modelling (ABM) approach combined with cognitive modelling techniques to give resultant reactions for each individual in the threat environment. A difficulty arises with current cognitive models such as SOAR (Laird, Newell, &



**Figure 1:** A simple four layer cognitive model for Real Time or large population agent-based modelling

Rosenbloom, 1987), PMFserv (Silverman, Johns, Cornwell, & O'Brien, 2006), or ACT-R (Anderson et al, 2004) that typically cannot adopt a spatial fight or flight behaviour set with respect to a three dimensional field affector in real time. Similarly, such cognitive models tend to view choice of behaviours non-spatially as a series of Performance Modifier Factors (PMFs) rather than being reflective of the evolutionary cognitive layers identified in brain structure studies (Leonard, 1993). For this reason we adopt a field-based partial subsumption architecture (Figure 1) simplistically based on identified evolutionary layers of brain structure (Givens, 2009) as applied to problem solving as a test of the utility of the approach.

In the above model each layer may be viewed as a *topology* that may linearly interact with other layers to achieve a resultant action (Johnson & Ivancevic, 2013). An ingredient missing from the above description is the impetus to move forward against a prevailing threat pressure to retreat. Social pressure from both command structures and social cohesion, along with the degree of internalisation of those pressures cognitively (aka *morale*), provide that impetus. Local minima in resultant topologies may produce indecision and delay as the agent weighs the threat against building social pressure to conform to command imperatives. This produces a Lewinian-styled (Lewin, 1936) attract-avoid model upon which can then be applied to the trained tactical knowledge of the individual being modelled.

To bring these ideas together, a model of combat behaviour has been developed that creates a "danger map" for each agent based on threat from enemy lines of sight, perception of generically dangerous parts of the terrain and own-side fire. It incorporates social pressure as a means to achieve forward momentum and a simple set of tactical principles which each agent uses where appropriate. This achieves a tactical outcome that dynamically represents the effects of suppression, morale and tactical training in a manner that allows behavioural adaptation to new environmental and enemy contexts seamlessly. We comparatively test aspects of this model against the CAEn behaviour set to determine changes in scenario lethality and mission success times.

## 2. MODEL IMPLEMENTATION

The model was implemented in NetLogo (Wilensky, 1999), a multi-agent programmable modelling environment in which entities (known as agents) can be assigned rule sets which allow them to interact with other agents and the grid based environment. In our scenario, each agent represents an individual soldier moving and interacting in a simplified urban environment. The environment contains two types of terrain features: open areas and obstacles. Agents may move onto open grid locations but are blocked from moving into obstacles. Similarly, agents can see across open areas but line of sight is blocked by obstacles. The set of obstacles are arranged into groups of rectangular shaped collections to resemble buildings in an urban environment. The locations of these buildings are then randomly placed into a 100x100 grid where a single grid notionally represents one square metre (Figure 2 – the pink squares on the corners of buildings will be explained in Section 4).

The agents in the scenario are split into two teams of five, coloured red and blue. Red team's mission is to defend the blue area surrounding the blue flag while blue team's mission is to take over the same area by defeating all red agents. Red agents are initially placed in the area in the top region of the environment closer to the blue flag, while blue agents are always placed in the bottom region and must traverse up the environment to reach the flag. Both teams do not possess knowledge of each other's presence until visual contact is made. At that point, both teams will engage and attempt to destroy the other team and continue to move towards the flag and capture the area.

Agents are equipped with a low fidelity representation of a generic small arms weapon. All shots fired land randomly on a grid location within a two metre radius from its intended target. If the grid location of the landing position is out of line of sight from the shooter then the shot is assumed to have hit an obstacle and does not land at its projected location. Shots landing on a grid location inhabited by an enemy have a chance of killing the agent according to some kill probability ( $P_k = 1 - \text{distance to target}/100$ ). Killed agents no longer participate in the simulation and the remaining agents cannot detect them.

## 3. AGENT BEHAVIOUR DESCRIPTION

Each agent in the model consists of four major sub systems that control its overall behaviour in the environment. The four subsystems include: 1) the informational layer, 2) the decision rule sets, 3) the firing and movement rules and 4) a reactive suppression model that considers an agent's visceral reactions to immediate threats.

### 3.1. Informational Layers

The informational layers are a set of influence maps (Tozour, 2001) overlaid over each grid location in the environment that detail certain characteristics about that particular location. These maps help to drive the decision and movement rules that are contained within each agent. The information contained in these maps may be used in isolation for particular movement rule sets or they may be aggregated to form a danger map representing a combined fused picture of the environment. The danger map is made up of four different influence maps; the terrain map, the incoming fires map, the enemy line of sight map and the friendly line of sight map.




Figure 2: Simplified urban environment.

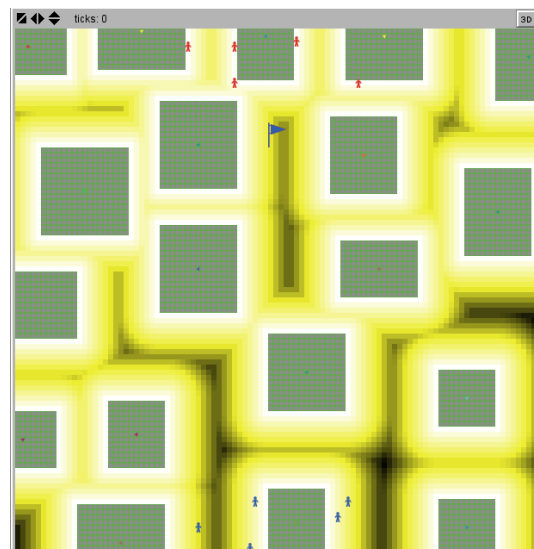


Figure 3: The terrain map.

The terrain map represents the perceived dangerousness of traversing over certain grid locations in relation to obstacles in the environment. Light coloured areas represent less dangerous areas than dark regions as shown in Figure 3. Generally, grid locations that are next to obstacle grids are safer than grid locations in more open areas because there are many more opportunities for an enemy to achieve a clear line of sight to a target standing on open ground than one that is standing next to an obstacle. We have not modelled building or terrain damage in this scenario so the terrain map remains static throughout the simulation.

The remaining three influence maps are dynamic in nature and are continually updated based on temporal events in the simulation. An example of such an event might be an agent suddenly appearing and placing fires on a particular location in the environment. The level of danger on the corresponding location (and neighbouring locations) on the incoming fires map is then increased. The increased danger levels of these grids then fade over time (if there are no further fires detected) to represent some persistent memory on behalf of the agent.

Knowledge about an enemy's location provides another layer of an agent's overall danger map via the enemy line of sight map (Figure 4). Danger levels are assigned to an agent's enemy line of sight map based on areas in the environment that an agent's known enemies can currently see from their present locations. The danger level of a grid location is directly proportional to its distance to an enemy and also proportional to the number of enemies that have line of sight to it. Agents periodically scan the environment to try and detect enemy agents and update their enemy line of sight map. The model assumes that agents on the same team have shared situation awareness so an enemy detection by one agent will be revealed instantaneously to all other agents on the same team. This danger map is dynamic and changes every time an enemy agent is detected moving within the environment. Similarly, the friendly line of sight map considers the line of sight from other agents on the same team so that agents avoid collateral fire when traversing the terrain.

Together, the combinations of the danger levels of the terrain, enemy line of sight, friendly line of sight and fires maps together form a fused representation of the danger or risk to an agent traversing across that environment. Agents then try to minimise their risk when moving across the terrain by only moving across less dangerous grid locations.

### 3.2. Decision rule sets

Each agent in the simulation is controlled by a state machine that determines what its current set of decision rules should be at any point in time (Figure 5). Events (both internal and external to the agent) generated will transition the agent to the appropriate states in the state machine. Blue agents begin the simulation in the patrol state in which they will move towards any assigned waypoints (in this case the goal location) until they come into contact with an enemy agent. At that point they will move to the nearest available covered region and proceed to engage the enemy. Red agents begin in the engage state and are ready for the defence of the blue flag. When combat initiates between the two teams, the battle will last until one side has been completely destroyed.




Figure 4: Blue team's enemy line of sight

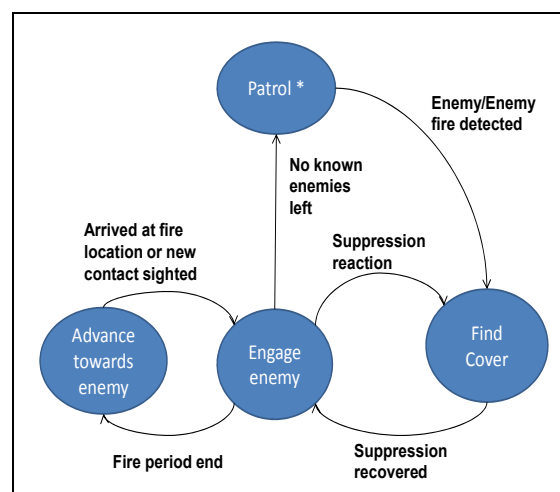


Figure 5: A soldier agent's internal state machine

### 3.3. Firing and Movement Rules

#### *Detection and Targeting algorithm*

A simplified detection algorithm was used where each agent automatically detects agents from the opposing team if it has line of sight to it during that time step. The enemy's locational information is also added to a dynamic map known as the *spot* map. The spot map is used as a memory persistent mechanism to record previous enemy detections. This is useful in the instances when the enemy has been spotted during one time step at a particular location but then goes out of visual detection in the next time step. Agents may then use any information present in the spot map to select a target area to orient the team towards.

Each team has an agent that is designated as the leader. The leader uses the information present in the spot map to designate a target area for the team to orient and move towards. The target area chosen is the centroid of the three most recent spot locations. If there are no recent enemy sightings, the goal flag will be designated as the target area. The selection of the actual target to engage is dependent on the firing location, as discussed below. If the current leader is killed at any point in the simulation, a new leader is appointed to continue the targeting process.

While in the engage state agents will fire for a random amount of time before deciding to move to a firing location that is closer to the target location. An agent can only move up to its next fire position if it is the furthest agent from the designated target and only if there is at least one other agent on the same team in the engage state to provide covering fire. This has the effect of moving the team towards the target area in a leap frog manner which forms the basis of the concept of fire and move.

#### *Selecting a firing location*

A set of possible firing locations is pre-calculated for each simulation based on areas in the environment which provide agents with opportunities to achieve line of sight to their targets but quickly move into cover if the enemy fires back. For the case of rectangular shaped buildings these locations correspond to locations adjacent to the corners of each building (shaded pink in Figure 2). When an agent reaches its selected firing position, it will attempt to shoot at the closest enemy within line of sight from its current location. It may also move a small distance (up to 4 grid squares) away from its firing location to achieve line of sight to its target.

The firing position is selected by first passing each firing location through a set of filters to help eliminate undesirable positions. These filters consider issues such as an agent's maximum bound length, the desired movement direction, occupation by other friendly agents and proximity to locations with line of sight to the target point. In the event that no firing location meets the strict filtering process, we relax the filters by selecting safe locations that moves an agent closer to the target point.

The firing locations remaining after filtering process are then evaluated based on the risk associated with moving towards that location. The path cost to each firing location is calculated via an A\* algorithm that finds the least dangerous route from an agent's current position to the selected position using information from the combined danger map. The firing location with the lowest path cost is then selected as the agent's next firing position.

### 3.4. Reactive suppression model

The previous section describes the agent's decision rule set in the combat cycle. However, each agent also contains mechanisms of a reactive suppressive system that supplants the combat decision rules at particular points in the simulation. The suppression model functions similarly to the one used by the Close Action Environment (CAEn) wargaming and simulation tool. Each shot fired by an enemy agent that lands within a user defined distance from its target increments an agent's suppression level (by 4 as default). In the absence of further enemy fire the suppression value will decay at a constant rate (by 2 each second as default). If the suppression value exceeds a user defined percentage of its maximum suppression level (default max = 10), then the agent has the same percentage chance of becoming suppressed. A suppressed agent will either drop to the ground if there is no immediate cover available or move to the nearest location out of line of sight to the enemy. The suppression increments caused by shots landing near agents that are either lying on the ground or behind cover is reduced to 10% of the default value. This provides an opportunity for agents to recover from suppressive fire if they are in areas or positions that provide increased protection. When the tank reaches a lower threshold (currently 10% of its maximum suppression level) the agent will transition back to the engage state and will no longer be considered as suppressed.

#### 4. MODEL COMPARISONS

Using the scenario described above, a small experiment was conducted to explore differences in agent behaviour and performance when various components of the above framework were added. In particular, we wanted to determine if improvements could be achieved when compared to a default CAEn agent. In order to do this, a simplified model of a CAEn agent was constructed that displayed only three basic behaviours: 1) path following towards the flag, 2) firing reaction upon visual line of sight contact with the enemy and 3) dropping to ground when suppressed. The CAEn agent served as our baseline agent to which additional behaviours were incrementally added until the agent displayed the full fire and move (F&M) behaviour set

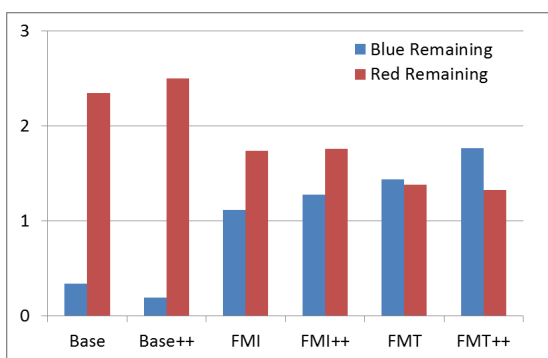
**Table 1.** The list of agent types that blue team adopts

Agent Type	Movement behaviour	Suppression behaviour	Wait for covering fire?
Base	Towards goal only	Drop to ground only	No
Base ++	Towards goal only	Drop to ground and move to cover	No
F&M Individual	Finds firing locations	Drop to ground only	No
F&M Individual ++	Finds firing locations	Drop to ground and move to cover	No
F&M Team	Finds firing locations	Drop to ground only	Yes
F&M Team ++	Finds firing locations	Drop to ground and move to cover	Yes


described in the previous section. The behavioural templates adjusted an agent’s movement behaviour, suppression behaviour and whether it takes advantage of covering fire. This created a total of six different behavioural types (Table 1) which were assigned to the blue team only. The red team always used the F&M Team++ option. Each scenario was run 100 times to produce a sample for statistical testing. Three metrics were recorded at the end of each replication: time taken for one side to win the battle, the final number of agents remaining for each team at the end of the run and the winning team. A win was recorded when there were no agents remaining on the opposing side.

#### 5. RESULTS

The results collected from the experiment described in the previous section show that the blue agents performed better when they were progressively endowed with additional information and behavioural capabilities. The number of blue agents remaining at the end of each run increases as blue behaviours become increasingly sophisticated (Figure 6). In addition, the number of red agents remaining decreases as blue adopts the more sophisticated behaviours. These findings are re-enforced in Figure 7 which shows the percentage of runs were blue was successful (i.e. there were no remaining red agents). As expected, all fire and move algorithms outperformed the baseline behaviours while team fire and move algorithms (FMT and FMT++) performed the best overall. However, there was no statistically significant<sup>1</sup> difference between any of the move to cover (++) behaviour types and their corresponding baseline behaviour of dropping to cover only. This can likely be attributed to the fidelity of the weapon representation in this model and requires further testing. In particular, a higher fidelity representation of the weapon and associated  $P_k$  values would almost certainly influence the outcome, particularly in relation to  $P_k$  values associated with moving targets. It should also be noted that in the FMT++ there is no statistical difference in the success rate of either team which is an expected result given that they both use identical algorithms in that particular case.



**Figure 6:** Average survivor count by team and behaviour type.



**Figure 7:** Blue success rate for 100 replications. Dashed vertical lines separate results that are statistically significant.

<sup>1</sup> A Shapiro-Wilk test indicated that the data was not normally distributed. Hence, a Kruskal-Wallis test was used to determine if data points were likely to come from the same distribution. Alpha values of 0.05 were used for both tests.

Scenarios where blue adopted the fire and move behaviour types also resulted in longer median times for the winning team to complete the mission (Figure 8). Of the fire and move behaviour types the ‘move to cover’ options produced statistically longer median times than the ‘drop to ground only’ options. This result combined with visual inspections of the simulation suggest that the more advanced behaviours produce the type of ebb and flow of combat that is not well represented in current simulations. The fire and move algorithms often produced cases where either team would gain some ground and then retreat or wait in cover for a short period before advancing again. In addition, a military subject matter expert suggested that the fire and move behaviours were more realistic than the baseline behaviours and that the military tactics observed were plausible given the assumptions and limitations of the model.

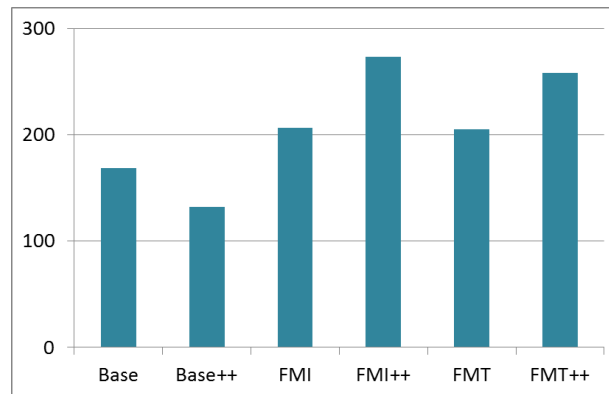


Figure 8: Median time taken to complete scenario by blue behaviour

## 6. CONCLUSION

This paper has pointed out the limitations of some of the human behaviours represented in existing combat simulations and suggested some additional simple dynamic behaviours built from a topological evolution-based framework to make them more representative of military tactics. A small experiment was conducted in order to seek quantifiable improvements in agent behaviour. The results indicate that the dynamic fire and move behaviour is an improvement over the base behaviour and produces plausible military behaviours, using a Lewinian threat field balanced with command social pressures exposing suppression, morale, and training to examination. However, we admit that this comparative study is by no means exhaustive. Further work is planned to test the behaviours over a wide range of terrains, force sizes, team rules and suppression variable combinations. Additional tests are also planned to measure the plausibility of the dynamic behaviours including a more rigorous face validation process and comparisons with data from field exercises.

## REFERENCES

- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). *An integrated theory of the mind*. Psychological Review 111, (4). 1036-1060.
- Givens, D.B (2009), Amphibian Brain, <http://center-for-nonverbal-studies.org/12801.html>
- Johnson, W.T. & Ivancevic, V.G. *Individual Psychological Conflict within Group Dynamics* MODSIM 2013
- Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). *SOAR:an architecture for general intelligence*. Artificial Intelligence, 1-64.
- Leonard C.M., et al. *Anomalous cerebral structure in dyslexia revealed with magnetic resonance imaging*. Archives of Neurology. 1993;50:461–469.
- Lewin, K. (1936). *Principles of Topological Psychology*. New York: McGraw-Hill Book Company.
- Shine, D. R. and Coutts, A. W. (2006). *Establishing a Historical Baseline for Close Combat Studies – The Battle of Binh Ba*. Proceedings of the 2006 Land Warfare Conference.
- Silverman, B. G., Johns, M., Cornwell, J., & O'Brien, K. (2006). *Human Behaviour models for Agents in Simulators and Games:part 1:enabling science with PMFserv*. Teleoperators and Virtual Environments
- Tozour, P. (2001). *Influence Mapping*. In M. Deloura, Game Programming Gems 2 (pp. 287-297). Charles River Media.
- Wilensky, U. (1999). NetLogo. <http://ccl.northwestern.edu/netlogo/>. Evanston, IL: Center for Connected Learning and Computer-Based Modeling, Northwestern University.
- Winter, R. P., & Clovis, E. R. (1973). Relationship of Supporting Weapon Systems Performance Characteristics to Suppression of Individuals and Small Units, TR 73/002. Sunnyvale, California: Defense Sciences Laboratories, Mellonics Systems Development Division, Litton Systems, Inc.