War Gaming Supported by Cognitive Computing and Time Manipulation

Doug Stapleton
IBM Australia Limited
28 Sydney Avenue, Forrest ACT 2603
AUSTRALIA

dougstap@au1.ibm.com

ABSTRACT

Cognitive computing is a disruptive new approach to developing war games and training forces. This approach can help overcome the current difficulty in scaling the translation of force level doctrine into a high level strategy game. Cognitive computing can generate likely scenarios and cascading effects that are based on rules generated from multiple sources of doctrine such as books, journals and after action reports. Adversary behaviour can be modelled from how they adapt to space and time conditions, current events and the positioning of friendly troops and military assets. Historical data can be synthesised into guidelines such as potential adversary range of movement or speed of movement. Players can then be afforded the ability to manipulate the fourth dimension of time. This allows players to optimise strategy collectively in ways that maximise resilience. Time manipulation allows users to explore doctrine and its implications within different scenarios and how those decisions cascade to other players. This paper explores the enhancements possible to traditional war gaming through the use of cognitive computing and time manipulation.

1.0 THE PROBLEM DESCRIPTION

A realistic war game is one in which the Opposing Force (OPFOR) behaves consistently with a real world opposing force. How is this achieved? When humans role play the OPFOR and implement their view of the strategy that would be adopted by an OPFOR in the real world, we are inherently limited by our training, experience and perspective. There is a natural tendency to make decisions the way that we would make them based on our culture and doctrine, not on those of a potential adversary. The thinking can effectively become Blue on Blue and degrade the training value of the war game. How can this limitation be overcome. We need to have a deterministic way of getting into the mind of the OPFOR commanders and making the OPFOR behave realistically. This may be done through programming the rules by which the game play is executed for the OPFOR, or by providing strategic guidance to the human OPFOR players.

This paper sets out an approach to generating the rules to achieve realistic and consistent OPFOR behaviour.

In discussing the representation of adversary forces, the NATO Code of Best Practice (COBP) [15] comments that "Historically, adversary capabilities and behaviours were often fully scripted or heavily constrained. This was more appropriate in Cold War contexts than it is today. However, it was never ideal for Command and Control (C2) analysis because the dynamic interaction among friendly, adversary, and other forces is a critical element of C2 representation. Today, much more robust adversary representation of operational capabilities and choices are employed and indeed are necessary. Analysts must consider not only a range of scenarios, but also the range of possible adversary actions and reactions." Time is an additional training dimension that analysts need to consider.

In the real world, time is a one way vector. However, there are occasions in strategy games where we would benefit from being able to reset the clock and take a different path forward. This paper also looks at how strategy games can be manipulated in time and the associated learning implications.

Indeed the world has become more complex and less certain, as Squadron Leader McPherson [16] puts it "Threats have become less structured, more complex and unpredictable whilst global terrorist organisations such as ISIS in Iraq and Syria, Al Qaeda, the Shabab in Somalia, and Boko Haram in Nigeria, have all emerged to dominate the headlines. Operations in Iraq and Afghanistan reinforced the need for a comprehensive and multi-dimensional solution to future NATO missions. Moreover, the emerging considerations of possible operations within NATO Member countries and the associated sensitivities this creates have only added to the dilemmas and challenges for commanders." Those non-state actors typically do not behave as traditional forces having more random elements to their behaviour profile, increasing the risk when only training with Blue on Blue doctrines.

1.1 Cognitive Computing

The term OPFOR in some ways is a diplomatic nicety to avoid the ramifications of naming a real nation as a likely enemy. However, ideally the OPFOR should accurately emulate real-life enemies so that their behaviour in tactics, techniques and procedures (TTPs) in war games and exercises has similar military characteristics to the expected real-world foes.

Cognitive computing is a way of mining a body or corpus of knowledge to derive behaviours. We don't want to be guessing at how an enemy commander would behave, we want to understand their military doctrine, their culture, their personality, in an organised fashion. This approach allows us to model the OPFOR on various real world scenarios. What is needed is a body of evidence such as:

- History of the OPFOR nation and the various military conflicts over time
- Their military doctrine as they define it and as we understand it
- After action and military intelligence reports that give factual evidence on how their military has behaved
- Specific information about their commanders, personality, bias for action and other factors that influence their military decision making
- General cultural proclivities of the OPFOR nation and how that might bias certain decision making processes

While the OPFOR can be modelled as an amorphous blob, it is more productive to be able to configure each OPFOR instance based on a rich knowledge of potential adversaries, with specific realistic behaviours providing a deeper understanding behind the decision framework/thinking that drives their doctrines.

1.2 Time Manipulation

Once a multi-player war game is played past a certain point, there is a cost involved in replaying the game to incorporate learning that would change earlier decisions.

Course of Action analysis does not always result in a single outcome that is clearly advantageous. The game may come to a certain point where the players would like to explore a number of options. Having the 'state' of the game saved after each cycle of movement allows the game controllers to reset the play to a certain point in time and then evaluate the various outcomes.

Thus time manipulation allows users to explore doctrine and its implications within different gaming scenarios and how those decisions cascade to other players. There is a broad topic of research called 'temporal difference', where learning outcomes have to be considered in context to the stream of events in which they occur. Hence time is a fundamental resource for learning.

1.3 Bringing it together

In the last hundred years, we have witnessed three distinct evolutions in warfare.

- In World War I, some battle fields were industrialised with light railways [20] to bring in men and materiel [21], with little actual territory changing hands but many thousands of casualties.
- In World War II, the war ending technology of the atomic bomb did not exist in an operational form at the beginning of the conflict. However the pace of the conflict was slow enough that technology continued to develop over the six years of the war. In that sense, WWII started to demonstrate the feedback loop of innovation, as the research and manufacturing necessary to operationalise the new atomic technology was integrated into national war planning.
- In today's world, the length of modern supply chains and the shortening timeframe of modern warfare means that we can no longer research and develop new weapons platforms, nor indeed re-order existing complex weapons platforms for delivery; they must be available in the necessary quantities before conflict arises. In the author's personal opinion, we will stand and fight; win or lose, with the resources at our disposal at the start of the next major conflict against a peer adversary.

This change and tightening in our ability to respond makes it all the more important to be confident that we are training for, and able to respond to realistic scenarios.

There is a fundamental difference when the game has its scenarios and behaviour modelled accurately on the opposing force doctrine and allowing for the game to be replayed from a certain point. This provides for a failure to be replayed and alternative courses of action experimented with to determine a winning strategy. The key point being that replays or temporal learning moves beyond teaching just the rules of the game and reinforces the strategy and principles to apply.

2.0 COGNITIVE COMPUTING

The term *cognitive computing* is describing a computing platform that is able to form concepts, understand and reason, and learn with the capacity to recognise patterns and use natural language to communicate.

That is not to suggest that computers have developed to the point of having human like intelligence characterised by perception, consciousness, self-awareness, and volition.

At a practical level, we might recognise intelligence in an OPFOR that generally acts in a recognisable way consistent with the opposing force military doctrine, culture and experience.

The following sections provide a background to the theory behind cognitive computing, the range of sources from which OPFOR behaviour can be derived and how that contributes to the modelling of adversary behaviour.

2.1 Cognitive Computing Theory

A military researcher might pose the question "What is the maximum speed of the OPFOR Main Battle Tank?" while other more complex concepts such as "Send out a reconnaissance in force" would be interpreted differently based on the applicable military doctrine.

We need computers to interact and reason over natural-language content in the same way that humans do. That development has occurred over a long period of time and has primarily advanced through techniques of open-domain question-answering (QA) [1, 2, 3].

To have computers reason like a human has required advances in many areas of computer science and artificial intelligence (AI), including information retrieval (IR), natural-language processing (NLP), knowledge representation and reasoning (KR&R), machine learning, and human-computer interfaces (HCIs) [4].

Essentially there is no one computer program that is able to reproduce the subtleties of the human mind and the way it understands language interactions. There are multiple analytical paths, each of which contributes part of the solution.

The challenge in computerising these multiple analytic paths was addressed by an architecture framework for integrating diverse collections of text, speech and image analytics called Unstructured Information Management Architecture (UIMA) [5] developed by IBM and later contributed to the Apache foundation [6] and is in use by industry and academia today.

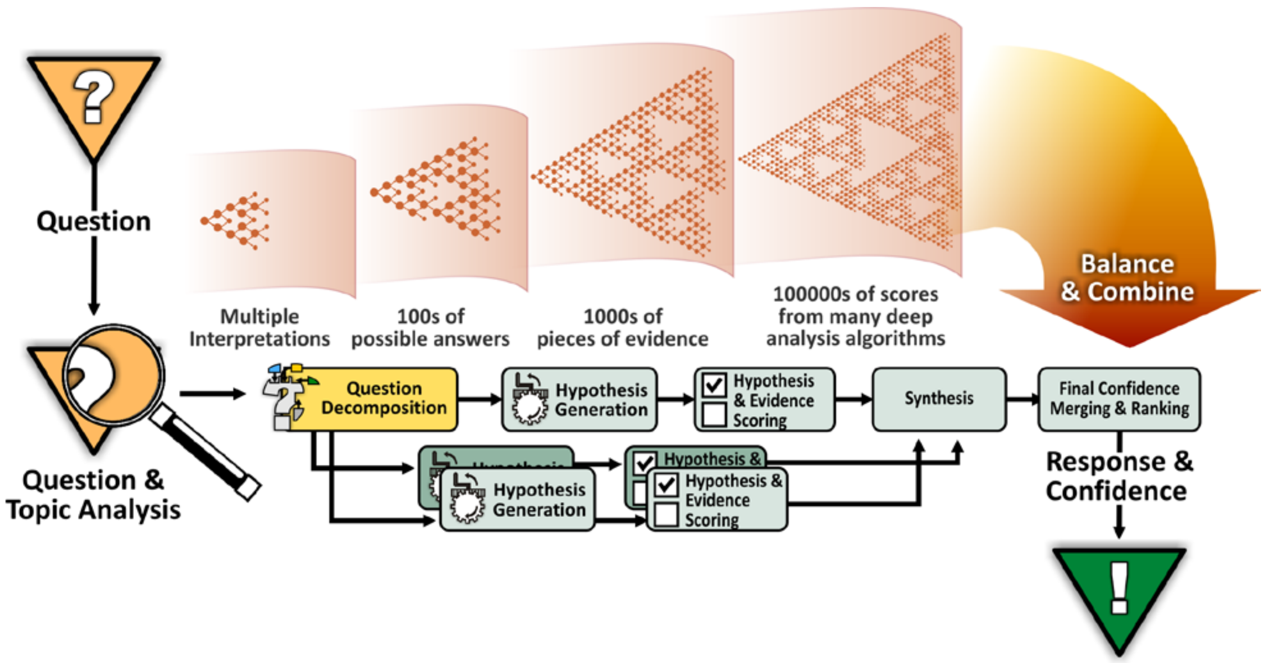

Figure 1: Question and Answer processing pipeline

Building on the UIMA architecture, a set of parallel processing pipelines is used to analyse the question and independently pursue possible candidate answers by searching many different resources. Evidence is then gathered for each alternative answer until a final weighting gives a confidence score for presentation of the preferred answer as illustrated in Figure 1.

A feedback loop where the system is trained by experts in their field improves the correctness of the answering process. Context is a key part of that training process; for example a reference to 'Atlas' would need disambiguation by context to either a member of the Atlas rocket family; a mountain range which stretches across north-western Africa extending about 2,500 km (1,600 mi) through Algeria, Morocco and Tunisia; or a collection of geographic maps. A key advantage provided by that expert training is making the results available to all subsequent queries. Future users interacting with the system will benefit from that expert training, which the system will 'remember' and apply to future questions.

Understanding the question using natural language parsing techniques is the first part of the processing pipeline. The class of thing being asked for is referred to as the *Lexical Answer Type* or LAT [7, 8]. Questions can be further decomposed into sub questions that can be independently answered [9].

The final stage in the pipeline architecture is to bring together the evidence for possible answers and to produce a ranking of the likelihood or confidence in a particular answer being the "correct" one with machine learning assisting in improving the overall ability of the system to rank correct answers [10].

To be practical, we need the performance of cognitive computing to be reasonable. An interesting example is when IBM built a computer system called Watson to compete on the U.S. game show Jeopardy! using the cognitive computing processes outlined above, the first computer attempts took two hours to answer each single question. Although correctly answered, it was hardly a champion performance level. The solution as part of the UIMA architecture was to introduce massive parallelism in the computations [11]. Ultimately 2,880 processors working in parallel on a single question brought the response time down under 3 seconds allowing the Watson computer to challenge and finally win Jeopardy! in 2011 against the two highest ranked human players [12].

Cognitive computing applied to the medical field has moved towards the concept of evidence profiles and supporting evidence that enables the user to explore the specific evidence justifying a diagnostic hypothesis or treatment alternatives [13]. In that approach, dimensions of evidence include clinical findings, demographics, symptoms, and family history, which contribute to the overall confidence score.

Cognitive computing applied to military research may also benefit from the use of evidence profiles that take advantage of specific military concepts and doctrine to evaluate alternatives for an OPFOR decision in setting a Course of Action, as an example.

2.2 Sources

Broadly speaking, four things are needed to model an OPFOR [17] as illustrated in Figure 2:

- the military platforms to be used
- the organisational structure of the armed forces and C2
- the commanders who will exercise force through the organisational structures available to them
- the dogma or doctrine driving the decision making processes

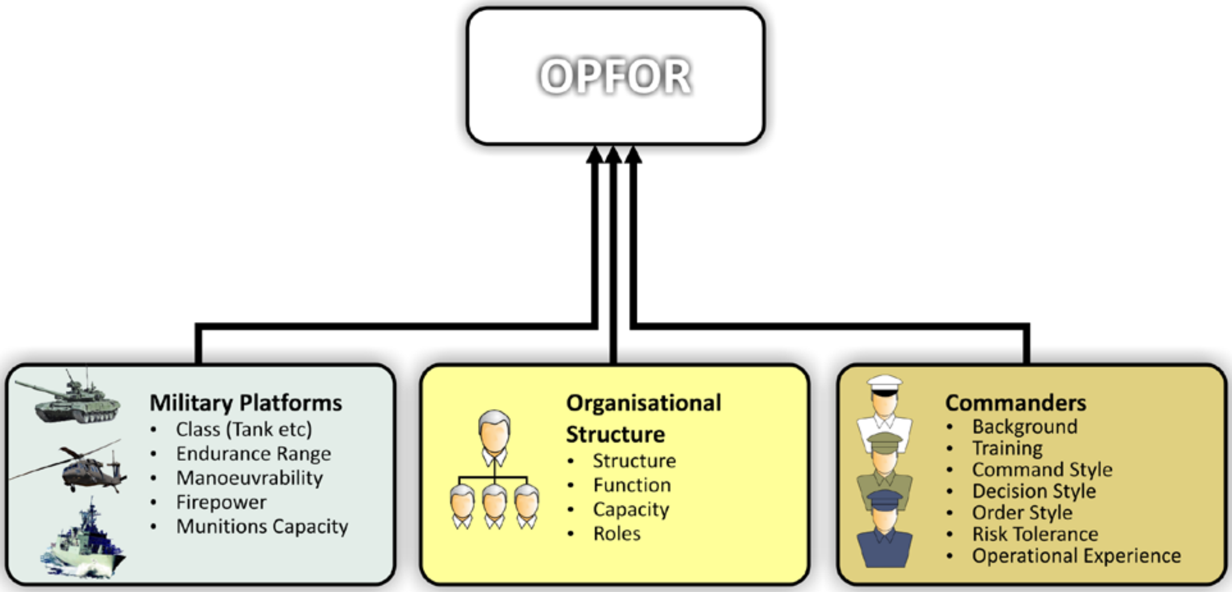

Figure 2: Modelling an OPFOR

Military Platforms - the characteristics of military platforms include their class (tank etc.), endurance range, manoeuvrability, firepower, munitions capacity.

Organisational structure - the principal differences between military organisations are related to structure, function, and capacity:

- Structural differences may relate to the number of layers, span of control and linkages between nodes.
- Functional differences arise from the distribution of responsibility, authority and information.
- Capacity issues may relate to personnel, communications and information, field training and operational experiences.
- Roles may be a useful way of summarising sets of behaviour. Expectations of key roles such as the military commander have evolved over time to be relatively well understood sets of behaviours.

Commanders - the attributes of commanders which can be modelled include:

- background
- training
- command style
- decision style
- order style,
- risk tolerance
- operational experience

Dogma or doctrine - all professional military organisations are guided by a recognised and defined military doctrine, or "way of doing things". That military doctrine is guided by the political background of the nation state and cultural influences that provide definite guiding viewpoints on such issues as minimisation or tolerance for civilian casualties. Non-state actors are more typically guided by dogma that may be religious in nature. It becomes much more difficult to find a published source of behaviours that non-state actors would consistently conform to. Cognitive computing is able to broadly mine a large corpus of social media and newsfeeds for sentiment analysis to form an aggregate picture of likely behaviours. Non-state actor's TTPs are perhaps best understood by mapping recent social media activity loci rather than measurement of their equipment technical capability.

In a modelling sense, even a generic OPFOR will have general characteristics. Military platforms move and fire weapons, commanders receive information and make decisions according to their organisational structure and support. The first two categories (military platforms and organisation structures) are objective and can have models applied to them, the third (commanders profile) and fourth (dogma or doctrine) are subjective.

However it would be far more useful, as is done in practice, to model the OPFOR on a realistic opposing force. Cognitive computing can query and derive fact based answers about the characteristics of military platforms from the literature corpus and can infer a hypothesis about a commander’s likely behaviour and personality characteristics.

All cognitive computing platforms require training, and as a high level concept, this can be viewed as the system providing a ranked set of hypothetical answers [19]. The 'ground truth' is determined by a subject matter expert and the answer sets are so marked. Next time the same question is asked, the system will rank the 'correct' answer higher and come closer to the way that an expert would answer the question. In the field of military intelligence, this cognitive learning ability provides a way of capturing the expertise of personnel that have spent many years analysing various subjects. By marking their version of the 'ground truth', they contribute their years of experience to the corporate knowledge on OPFOR matters.

Published or classified documents on the life and battles of historical commanders provide a depth of insight into their behaviours, which can be built into a realistic OPFOR model. A country specific OPFOR model may have a general set of cultural characteristics, over which is layered the specific personality attributes of each individual commander (noting also the rise of non-state actors as potential military adversaries). In an ideal world, the simulation scenario would be able to select an OPFOR modelled on a specific country (or non-state military organisation) with contemporary commanders for the most realism.

Historical commanders have a wealth of published information, so that we can potentially go into battle against the great commanders of history such as Napoleon, Patton, Montgomery and Rommel. This also creates the intriguing possibility that we can make the commander more or less conservative; more or less willing to take risks etc.

2.3 Modelling Adversary Behaviour

Traditional information gathering techniques to accurately model an OPFOR are usually based on human led research that builds up a particular level of expertise in the analyst team, whether that relates to military platforms, organisational structure or commander biographies.

Cognitive computing can go a long way towards automating that process. The sheer ability to absorb a large corpus of documentation including biographies, intelligence reports, after action reports and engineering specifications distinguishes cognitive computing from the manual efforts of human researchers [18].

Let cognitive computing search the corpus of material and suggest hypotheses to answer the various questions. Have a subject matter expert review those hypotheses and mark the best answers as 'ground truth'. That way the corporate knowledge of those experts can be retained and made readily available to future researchers. The place of cognitive computing in this overall process is illustrated in Figure 3.

The approach suggested here uses cognitive computing to gather the data and hypotheses on which the adversary will act. These are validated by subject matter experts and used in the existing adversary behaviour engines.

The issue of scale now becomes apparent. Cognitive computing techniques allow for substantial automation in determining OPFOR characteristics. There is no longer the need to have an amorphous OPFOR due to a lack of human resources for research. Realistic OPFOR modelling can be introduced based on the characteristics of real countries and political entities as shown in Figure 4. It will become a realistic proposition to accurately model all the potential adversaries, so that military exercises can have the maximum realism.

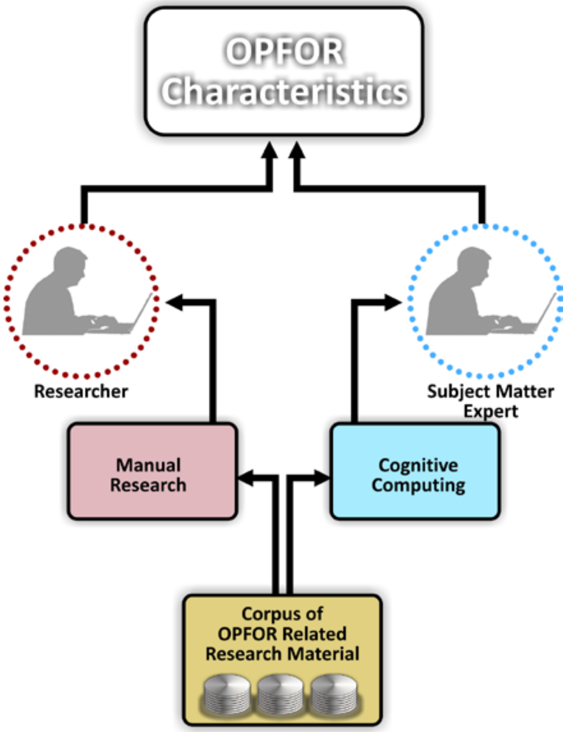

Figure 3: The place of cognitive computing in determining OPFOR characteristics

Figure 4: Modelling Realistic OPFOR

Envision a future where the military exercise planners can select a hypothetical but realistic opposing force that uses current commanders with the current organisational structures and available military platforms, based on real world scenarios.

3.0 TIME MANIPULATION

Players in a multiplayer real-time strategy game can be afforded the ability to manipulate the fourth dimension of time (4D). This allows players to optimise strategy collectively in ways that maximise resilience. Time manipulation allows users to explore doctrine and its implications within different scenarios and how those decisions cascade to other players.

3.1 The principles of time manipulation

We normally think of time as something that cannot be altered in the past. Given a particular scenario, we can only affect the future as we make decisions in the present. This also typically applies to game playing; our moves are based on the game up to now and new moves affect the play based on where the combined play is up to. It would be useful to be able to stop a chess game when we recognise an imminent check-mate if we continued on the present path, and back up to a safe position and branch off in a new direction.

A new class of 4D game engine provide the ability for players in a multiplayer environment to alter the position of the game on the overall timeline [14]. This requires the game engine to preserve the state of the complete set of game objects at defined points in time. A time slider as illustrated in Figure 5 allows the game to be reset to a previous safe rest point.

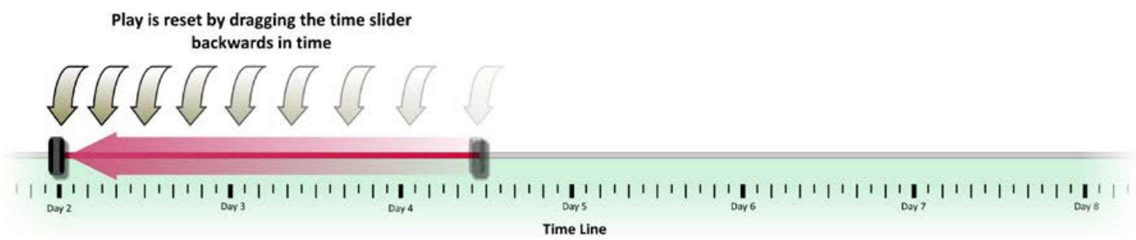

Figure 5: Time slider to manipulate the fourth dimension of time

As an illustration, let's assume that all military activity takes place during daylight hours only. There would then be a safe rest point at midnight where the position of all game objects would be consistent and able to be restarted from that point in time. This would give the game controller the option of restarting the game at an earlier day in the developing battle. Naturally more fine-grained control than this is needed. Smaller time units can be set, or on alternative moves by each side and so on.

The strategic point is that players can recognise when they are on the point of losing and collaborate with other players/instructors to optimise strategy collectively in a way that maximises resilience. This multiplayer time manipulation can help players explore how adversaries can thwart plans, create distractions and exploit weaknesses. It also provides a balanced viewpoint that can be crucial to winning the battle, such as the observable truth that armour columns must be supplied with both fuel and ammunition to be effective, and not get too far ahead of their supply columns.

4D time manipulation in multiplayer real-time strategy games allows players to:

- Optimise strategy collectively in a way that maximises resilience

- Explore doctrine and its implications within different gaming scenarios
- Affect change in different positions on the timeline, which creates effects that cascade to other players

3.2 The alternative theory of time

Somewhat like the old Irish joke "Well I wouldn't be starting from here"; the issue is that when considering time manipulation in a game, we need to think of how practical is it to be somewhere else and when would that decision need to be made?

As a simplified example to illustrate the concept within a larger scenario, let's assume that materiel is being brought to a port to supply the force. The ship departs from Port A on day four of the exercise and has 1000nmi to travel to reach Port B where it will disembark the materiel. Other players on day six of the exercise call out for the ship as they now urgently need those supplies. Assuming a maximum speed of 250nmi per day, the ship has to be somewhere within the circle of maximum range, as illustrated in Figure 6.

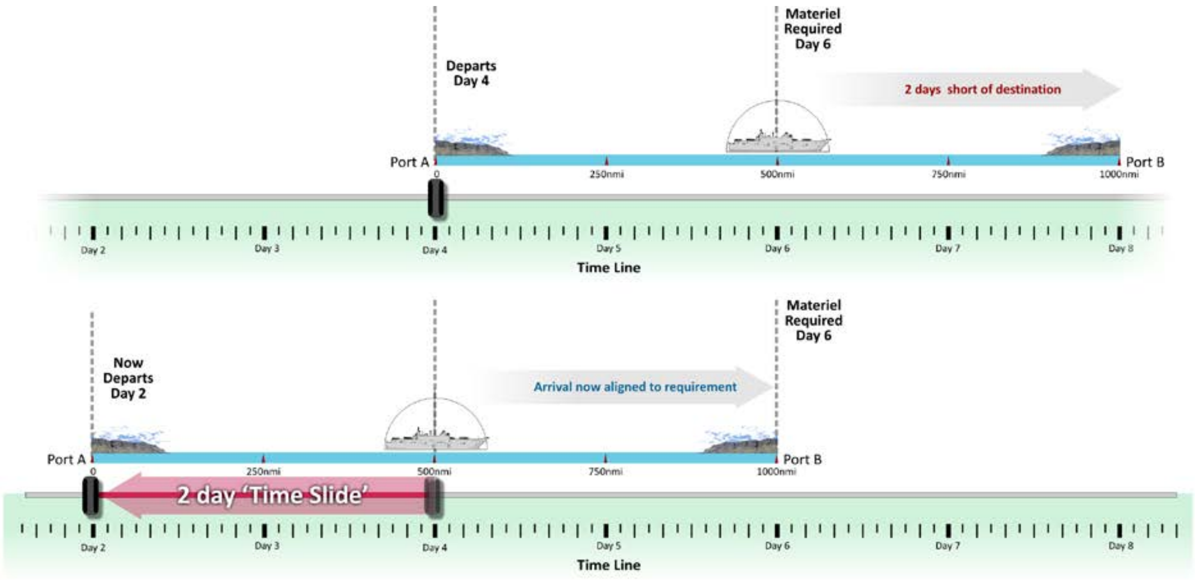

Figure 6: Illustration of ship and circle of maximum range that increases with time

Clearly the ship still has 500nmi or two days travel left. In simple terms, if the ship is out of position by 500nmi and can steam at 250nmi per day, then it is clear that the decision had to be made two days earlier than the ship actually left, to get the ship into position at Port B on time. The learning from this illustration is that the ship had to leave on day two of the exercise. By agreement, the time slider can be reset at day two and game play resumed with the ship being sent off on day two of the exercise to arrive in port on day six as required. There are of course other complex factors that all interact and would need to be thrashed out into the overall strategy. The ability to work this through as game theory can provide tangible feedback into doctrine, tactics, techniques and procedures.

3.3 Game play rules of time

By going backwards in time and letting the game play branch off into a new version of reality, the position of all the key objects must be known and the alternative decisions are implemented from that point. This means that objects cannot be moved arbitrarily around, they must move according to their established rules of behaviour; the ship can still only move 250nmi per day, regardless of when it is set in motion.

We can see that there is some skill involved in working out when to reset the time stage of the game to allow for the alternative directions to achieve the desired result. If the ship is needed 500nmi away from where it is at the point that the first play is halted, then clearly to obey the rules (such as maximum transit per day of 250nmi), there is a calculation involved which shows that the ship must have departed at least two days prior to when it originally departed. That new behaviour may also have other implications such as loading with cargo, assembling the crew, and so on, depending on the realism of the game play.

3.4 Independent time manipulation

Some game engines now allow for players to independently decide on where in the game timeline they want to play. So as a scenario develops, sensing imminent defeat, a player can independently move their time slider backwards to go back and marshal greater battle field resources. This is then propagated forward as a 'time wave' so the player reappears in the present with greater resources with which to deter an impending attack.

While an interesting conceptual development, this sort of time manipulation is perhaps more of a recreational game playing option, than for serious war game development.

4.0 OPTIMISING STRATEGY

We can now bring the two aspects together of cognitive computing and time manipulation and see the improvement in war gaming that this makes possible.

Realistic modelling of the behaviour of an opposing force is the key to building a credible training experience. Using cognitive computing to draw out the behaviours of the OPFOR and assist the rule creation for the adversary behaviour modelling will improve the realism of the decision making of the OPFOR.

As a war game is played out there may be a certain point where the game controllers deliberately want to branch out in various directions and see how the different approaches work towards the overall result. This deliberate choice implies that the game can be halted when a clear result can be foreseen; "continuing on this path we are going to lose, let's go back and try the other path".

However, there is also the situation where game play continues and the players themselves realise that they should have done things differently, such as in our simplistic example of sending materiel by ship two days earlier. The ability to manipulate the time slider by common consent is a significant efficiency improvement in that game play can resume at a sensible point without the overhead of playing the game to its logical conclusion or restarting the game from the beginning.

These combined abilities of cognitive computing and time manipulation can have a significant impact on the training outcomes. Perhaps the cognitive computing has supplied rules to the OPFOR game engine that when confronted by a large scale tank attack the OPFOR commanders general tactic is to pull back quickly deep into their own territory. The blue force armour in pursuit may end up moving faster than fuel or ammunition can keep up. However, replaying that scenario may help the war game participants to move logistics supplies to keep up with the battle force. The realistic response of an appropriately trained OPFOR game engine can reinforce key learning outcomes that directly translate into real world war fighting.

In future conflicts, one of the first commands may be to move the ammunition.

5.0 CONCLUSION

This paper explored the enhancements possible to traditional war gaming through the use of cognitive computing and time manipulation.

Cognitive computing provides an organised way of extracting behavioural indications for the opposing force which can be structured as rules of OPFOR behaviour.

Time manipulation gives the opportunity to implement learning in a faster cycle by being able to restart the war game at various points according to the learned outcomes.

Considering the length of modern supply chains, when we do have to stand and fight against a peer adversary, cognitive computing and time manipulation may have enabled our war gaming to focus on the essential preparation of resources for the next major conflict.

6.0 REFERENCES

[1] R. F. Simmons, "Natural language question-answering systems:1969", *Commun. ACM*, vol. 13, no. 1, pp. 15–30, Jan. 1970.

[2] M. Maybury, *New Directions in Question-Answering*. Menlo Park, CA: AAAI Press, 2004.

[3] T. Strzalkowski and S. Harabagiu, *Advances in Open-Domain Question-Answering*. Berlin, Germany: Springer-Verlag, 2006.

[4] D. A. Ferrucci, Introduction to "This is Watson", *IBM J. Res. & Dev*., vol. 56, no. 3/4, Paper 1, pp. 1:1, May/Jul. 2012.

[5] D. Ferrucci and A. Lally, "Building an example application with the unstructured information management architecture," *IBM Syst. J.*, vol. 43, no. 3, pp. 455–475, Jul. 2004.

[6] Apache UIMA. [Online]. Available: http://uima.apache.org/

[7] A. Lally, J. M. Prager, M. C. McCord, B. K. Boguraev, S. Patwardhan, J. Fan, P. Fodor, and J. Chu-Carroll, "Question analysis: How Watson reads a clue," *IBM J. Res. & Dev*., vol. 56, no. 3/4, Paper 2, pp. 2:1–2:14, May/Jul. 2012.

[8] M. C. McCord, J. W. Murdock, and B. K. Boguraev, "Deep parsing in Watson," *IBM J. Res. & Dev*., vol. 56, no. 3/4, Paper 3, pp. 3:1–3:15, May/Jul. 2012.

[9] A. Kalyanpur, S. Patwardhan, B. K. Boguraev, A. Lally, and J. Chu-Carroll, "Fact-based question decomposition in DeepQA," *IBM J. Res. & Dev*., vol. 56, no. 3/4, Paper 13, pp. 13:1–13:11, May/Jul. 2012.

[10] D. C. Gondek, A. Lally, A. Kalyanpur, J. W. Murdock, P. Duboue, L. Zhang, Y. Pan, Z. M. Qiu, and C. Welty, "A framework for merging and ranking of answers in DeepQA," *IBM J. Res. & Dev*., vol. 56, no. 3/4, Paper 14, pp. 14:1–14:12, May/Jul. 2012.

[11] *UIMA Asynchronous Scaleout*, The Apache Software Foundation, Apache UIMA Development Community, ver. 2.3.1, 2011. [Online]. Available: http://uima.apache.org/d/uima-as-

2.3.1/uima_async_scaleout.html

[12] E. A. Epstein, M. I. Schor, B. S. Iyer, A. Lally, E. W. Brown, and J. Cwiklik, "Making Watson fast," *IBM J. Res. & Dev.*, vol. 56, no. 3/4, Paper 15, pp. 15:1–15:12, May/Jul. 2012.

[13] D. Ferrucci, A. Levas, S. Bagchi, D. Gondek, and E. T. Mueller, "Watson: Beyond Jeopardy!," *J. Artif. Intell.*, 2012.

[14] P. Boinodiris, L. Sokol, R. Behrens, H. Blejer, *Wargames Powered by IBM Watson: Multiplayer Real-Time Strategy in 4D*, REDP-5128

[15] *NATO code of best practice for command and control assessment*.—Rev. ed. 2002, ISBN 1-893723-09-7, page 200

[16] Squadron Leader Colin J. Macpherson, Royal Air Force "OPFOR Opposing Force - The Enemy From Within," *The Three Swords Magazine* 27/2014, pp. 29-34

[17] *NATO code of best practice for command and control assessment*.—Rev. ed. 2002, ISBN 1-893723-09-7, Chapter 6

[18] R. High, *The Era of Cognitive Systems: An Inside Look at IBM Watson and How it Works*, REDP-4955

[19] R. High, B. Rapp, *Transforming the Way Organizations Think with Cognitive Systems*, REDP-4961

[20] I.M. Brown, *British Logistics on the Western Front: 1914-1918,* page 193

[21] K.D. Stubbs, *Race to the Front: The Materiel Foundations of Coalition Strategy in the Great War* , page 176