9th International Command and Control Research and Technology Symposium

Modelling Decision Making to Support NetCentric Warfare Experimentation

Mark G. Hazen

Group Leader, Virtual Combat Systems
Defence R&D Canada - Atlantic
9 Grove St, Dartmouth, NS Canada
B2Y 3Z7
Tel. 902-426-3100 x176 / Fax 902-426-9654
Mark.hazen@drdc-rddc.gc.ca

Matthew Fewell

Maritime Operations Division
DSTO Systems Sciences Laboratory
PO Box 1500
Edinburgh, SA 5111 Australia
Matthew.Fewell@dsto.defence.gov.au

Abstract

Netcentric warfare and associated concepts have a fundamental assumption that improved information infrastructures will improve military decision-making and therefore military effectiveness. Previous work on linking NCW applications to military effectiveness using modelling and simulation have had difficulties in modelling the decision-making aspects of the process, and in particular modelling them such that NCW applications can be shown to make a difference. A survey of conceptualisations of decision making suggests that two characteristics of a decision, its speed (or timeliness) and its soundness, and two types of decision making, recognition-based and analytic, constitute a sufficient foundation for modelling. Models based upon empirical timings can be found for decision speed for the two types, however modelling decision soundness is more problematic. A theoretical structure for decision soundness modelling is proposed using situation and decision spaces and constrained by commanders intent. This structure appears to be rich enough in detail to encompass the concepts of commanders intent, decision-makers beliefs and biases etc., but has not yet been practically implemented. Within the structure, situational awareness is represented by a sub-space of possible situations and perturbed by decision-maker biases, commander's intent is represented by a second subspace of possible situations, and decisions are functions from one space to another. Decision soundness is a metric to assess the results of applying a decision function on both the perceived and actual situations and compares these results to the commanders intent or mission objective. The application of the proposed structure to the analysis of maritime interception operations is explored.

Introduction

Network-Centric warfare (NCW) and associated concepts have a fundamental assumption that improved information infrastructures will improve military decision-making and therefore military effectiveness. Previous work on linking NCW applications to military effectiveness using modelling and simulation [1] have had difficulties in modelling the decision-making aspects of the process, and in particular modelling them such that effects due to NCW applications can be measured.

At base, the technology enabler behind network-centric warfare is the network. The key source of advantage therefore lies in the sophisticated generation, communication and consumption of information. To realise the advantage, the consumption step must lead to better decision-making. Hence, models aiming to determine the magnitude of the effectiveness gain from the adoption of NCW must address in some way the issue of modelling decision-making. To support the study of weapon systems and military operations, models of human decision-making are required that do more than assume that information of complete accuracy and comprehensibility arrives at the decision-maker when it is needed, is processed optimally, and is implemented flawlessly are required. Previous work [2] indicated that it is not the quantity of data or indeed the individual datum quality, but a context-induced information quality that is important in determining value to a decision-maker. It is believed that Pigeau and McCann's [3] Commander's Intent provides this context.

The study of human performance in decision-making is a major and active branch of both cognitive science and ergonomics. However, the majority of the research is focused on determining specific aspects of cognition or behaviour. Important influences range from the environment, such as stress, heat and cold, presence of the enemy, through factors such as level of morale, training and experience, to a wide range of cognitive biases that are innate, in the sense that they are difficult to avoid (to say the least), even with explicit training. The list of biases includes some quite surprising effects, such as the resource-overuse bias and the possible deleterious effects of thinking too much, in addition to the more well-known biases due to perception, confidence and problem framing [4]. How much of this detail should be included in a model of military decision-making depends on the intended use of the model, but model mechanics ought at least to be consistent with what it is known about decision-making performance modulators and biases. For example, it should not be assumed that any given type of decision—e.g. analytical decisions—is deterministic. That is, decision-makers faced with the same situation need not come to the same decision, even when situational awareness is deeply shared. This is particularly true of decision-making under uncertainty, a very common feature of military decision-making.

The process of decision-making, viewed as a whole, is widely recognised as cyclic in essence. The Observe Orient Decide Act (OODA) loop, its variants and extensions provide a conceptualisation that is useful for modelling military decision-making, and have been used successfully for the study of network centric warfare. In terms of a scheme for classifying decisions the cognitively-based classification scheme—analytical vs. recognition-primed—appears useful for the purpose of modelling decision-making [4,5]. Working from the precept of not making a model more complicated than is needed for the purpose at hand, we propose that the quality of military decisions be modelled by specifying just two characteristics: speed and soundness [4]. Speed is defined in terms of the time taken to carry out certain parts of the OODA loop. Soundness is more problematic and has been assessed in some studies using subject-matter experts (SME). A major aim of this work is to propose a framework for obtaining more quantitative measures.

From the History of Decision Modelling

Attempts at mathematically conceptualising decision-making have a history stretching back several centuries. That the problem remains open is an indication of its difficulty. In surveying the history of models of decision-making, one can discern two broad themes that are not always clearly separated. On the one hand, some models of decision-making aim to assist decision-makers to make better decisions. That is, such work typically begins with a paradigm of best practice in decision-making—for example, that Bayes's theorem ought to be consistently employed—and produces a tool to help decision-makers to achieve the ideal. On the other hand, decision-making may be modelled with the aim of describing how decisions are actually made in practice and quantifying the quality of those decisions. We pursue the second aim in that we are interested in how military decision-makers actually function, and from that to learn about how NCW may affect the process. To this end, and in view of the problematic nature of quantifying soundness, we focus in the following brief historical review on concepts of decision soundness.

For most of the history of decision-making modelling, the first of the themes was foremost. Prior to about 1700, it was generally thought that optimum decision-making should be conceptualized in terms of *expected values*: the best decision or sequence of decisions is that which secures the highest (or lowest, as the case may be) average value of a quantity over the long term. In fact, humans are naturally more risk-averse in their decision-making than is implied by theories based on expected values. For example, most of us will pay in a lifetime more in insurance premiums than we receive in payouts. Therefore, according to expected-value theory, the purchase of insurance is a bad decision. This conclusion overlooks the purpose of insurance, which is to buffer the effects of rare catastrophic events. The insurer assumes some of the risk of such events.

In the early 1700s, Bernoulli proposed the replacement of expected value with *utility*, a construct intended to capture the decision-maker's attitude to risk [6]. The decision-maker then seeks to maximise expected utility. This theory was too *ad hoc* to be widely accepted until it was placed on an axiomatic basis by Von Neumann and Morgenstern in the early 1940s [7].

Expected utility may well serve as a basis for modelling ideal decision-making, but from the descriptive rather than the normative point of view—how we actually make decisions rather that how we ought to—the now well established cognitive biases that seem innate in humans do not always follow Von Neumann's and Morgenstern's axioms [8]. This has prompted the development of a range of variant theories in recent decades, with limited success. For example, Dorsey and Coovert recently published a formalism based on fuzzy logic [9] that is interesting, but seems of limited usefulness in modelling military decision-making.

Cognitive task analysis [10] constitutes a different approach. It attempts to map in detail the mental processes needed to make a decision in a given context. This can be a formidable task: the number of possibilities, branches and contingencies grows exponentially with the complexity of the decision context. Although cognitive task analysis can be viewed as producing a description of the decision space, the task of assessing the soundness of decision options is not an aim of the analysis. However, the need for a quantification of soundness is the main requirement of the modelling required to support NCW analysis.

The result is that at present, the modelling of decision-making and indeed human behaviour modelling remains open at the most basic level. Extension of methodologies to the quantification of decision soundness in general has required the use of subject-matter experts. In the next section, we propose an initial modelling framework as a contribution to the advancement of the subject area. This is followed by a section detailing some of the issues that will need to be addressed to obtain a practical implementation of the framework. Although a full implementation is not yet available, we apply the framework to a Maritime Interdiction Operation scenario that has previously been analyzed for NCW effects. [1]

Proposed Conceptual Framework

We seek a model that covers the orient and decide steps of the OODA loop. However, although the 'act' part of the OODA loop seems problematic from the modelling point of view, some representation of it and the 'observe' part are necessary to deal with situations involving a sequence of decisions. Since a primary purpose of the model is to enable the study of NCW effects on decision-making, it must take as input some representation of the decision-maker's level-1 situational awareness [11] and commander's intent [3]. Similarly the metric of effectiveness, decision quality, needs to be calculable from the output. Hence, the output should be the decision speed, decision soundness, and the decision itself.

Situation and Decision Spaces; Decision and Implementation Functions

Our concept of the overall structure of such a model is shown in Figure 1. Its essence is a mapping between the physico-temporal and cognitive worlds. The physico-temporal world is represented by the '*situation space*' *S*, which describes all possible situations that could be faced by the decision-maker; each point *s* in *S* corresponds to a distinct situation. Following the ideas of situational awareness, '*situation*' includes not only entity locations but also immediate entity aims. Thus, identical physical scenarios with differing entity courses of action are different situations.

Deleted: Figure 1

Inserted: Figure 1

Deleted: Figure 1

The cognitive world is represented by the '*decision space*' *D*, which describes all possible decisions. Decision-making and implementation are then functions between the two spaces. The form of the decision function, $F: S \rightarrow D$, must reflect the decision-maker's understanding and abilities, and the form of the implementation function, $\delta: D \otimes S \rightarrow S$, describes how a particular decision will change the current situation when it is implemented.

We do not specify the structure of *S* and *D* beyond referring to them as 'spaces'. It may appear that they have the character of vector spaces in the sense that the objects *s* and *d* are assumed to be categorisable in some appropriate fashion. However, the mathematical construct of the vector space requires the properties of vector addition, scaler multiplication, null vectors and closure, some of which are unlikely to be definable.

The combination of the two spaces and linking functions has the appearance of the classic OODA loop model of decision-making. It should be possible to set up the decision function *F* to represent 'orient' and 'decide', leaving 'act' and 'observe' to be described by the perhaps lower-fidelity implementation function δ. If such a separation can be achieved, then decision speed would be given by a time associated with *F*. Whether this time is an input to or an output of the

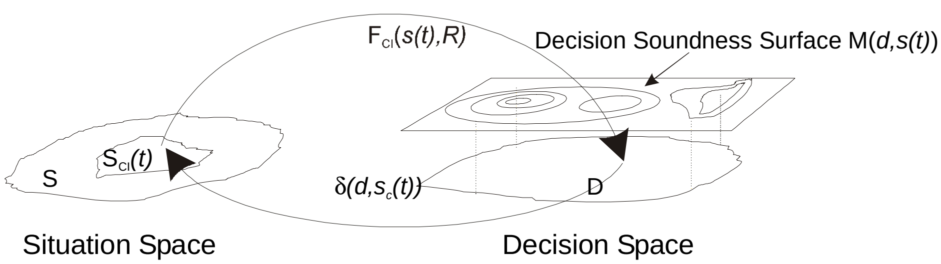

Figure 1: High-level view of the structure of a model of military decision-making. The 'situation space' S comprises all possible situations that the decision-maker may encounter and the 'decision space' D all possible decisions. The two are connected by a function $F_{CI}(s,R)$ *representing decision-making in response to the commander's intent and constraint set* R *and a function* $\delta(d,s)$ *representing the implementation of the decision. The soundness of the resulting decision is given by a function on* $D \otimes S$ *that returns the soundness surface* M.

model depends on how the model is instantiated.

Decision Function and Commander's Intent

The nature of the decision function F must reflect the decision-maker's creativity and freedom to generate options, the constraints on those options, and how to choose between them. There are therefore two main non-physical inputs to the function: the decision-maker's profile (experience and biases) and the mission context contained in the commander's intent. We define a set R of constraints on the decision function, which represents aspects such as the rules of engagement (ROEs), the decision-maker's authority to command resources, etc. The resulting decision function is labelled F_{CI} to emphasise the part played by the commander's intent in selecting the best decision for a given situation. Its physical arguments are the specific situation s at time t and the relational constraints are specified by R; its result is the decision d:

$$d = F_{CI}(s(t), R) \tag{1}$$

The *subspace* S_{CI} shown in Figure 1 indicates that usually the commander's intent does not encompass all possible situations. That is, while the domain of F_{CI} is all of S, for s not an element of S_{CI}, $F_{CI}(s,R)$ results in the decision to request changes to or seek clarification of the commander's intent. We make the following definitions:

Deleted: Figure 1

Inserted: Figure 1

Deleted: Figure 1

1) $s_C(t)$ is the true situation at time *t*.
2) $\hat{s}(t)$ as the situation at time *t* as the decision-maker understands it.
3) $S_g(t)$ is the goal of the commander's intent, which is likely to be a set of situations.

When $\hat{s} \neq s_c$, F_{CI} is applied to $\hat{s}$ (not s_c) to obtain a decision, but the implementation function δ is always applied to s_c, the true situation. Thus, if $\hat{s} \in S_{CI}$ and s_c is not, then a decision requiring substantive implementation is still generated, whereas when $\hat{s} \notin S_{CI}$ but s_c is, an unnecessary request for clarification of intent is generated. Practically, for a well constructed commander's intent, the probability of s_c not being in S_{CI} should be small.

The function $F_{CI}(s,t)$ connects the physico-temporal and decision spaces in a manner that depends on the commander's intent, but it is not the intent itself. Commander's intent, as defined in this paper, consists of the commander's full appreciation of the battlespace and how it will change over the operation. Thus, it consists of the current situation $s_c(t_0)$ when it was developed, the set of expected situations S_{CI}, the goal set S_g and the constraints R on the actions of subordinate commanders (i.e. rules of engagement and relationships between resources). When viewed over time, commander's intent can be visualised as specifying a corridor through S from $s_c(t_0)$ to S_g. In principle, it might specify a single path through S, but usually this would be regarded as micromanagement—the doctrine of 'mission command' posits that commander's intent should specify the goal without being too prescriptive on how to get there.

Structure of the Spaces; Situation Increments

In S, the 'sum' of two situations has no natural definition for any general pair of situations, but it is sufficient for our purposes to define a restricted change operator as follows. For any pair s_1, s_2 of situations, it is natural to conceive of a situation increment $\delta s_{1\to 2}$ as representing the 'difference' between the two situations; it indicates what is needed to change s_1 into s_2. Hence, it makes sense to write

$$s_2 = s_1 + \delta s_{1\to 2} \tag{2}$$

as a definition of the symbol '+' in its application to situation-like objects. Of more interest is viewing $\delta s_{1\to 2}$ as the result of a decision that changes s_1 into s_2. Thus (and with an obvious modification of the notation), $\delta(d,s_c(t))$ can be defined as a function from $D \otimes S$ to S such that

$$s_n = s_i + \delta(d,s_c(t)) \tag{3}$$

is the new situation following the implementation of decision d to situation $s_c(t)$ added to situation s_i. Developing this further, we see that $\delta(d,s_c(t))$ provides the feedback from the decision space to the situation space, as indicated in Figure 1.

Deleted: Figure 1
Inserted: Figure 1
Deleted: Figure 1

A major assumption of the model is that of closure for both D and S, that is, that *all* possible decisions are represented in D and every possible situation has a representation in S.

Decision Soundness

Soundness must ultimately be judged by reference to the effect of the decision on some measure of overall military effectiveness—absolute soundness—but usually this is not known to a decision-maker at the time when the decision must be made, so decision-makers judge the soundness of the available options by reference to the prevailing commander's intent as they understand it. This provides, we believe, a straightforward description of how military decision-makers operate and how less than optimal decision options (as measured on the absolute scale) come to be chosen. These considerations also indicate at least two aspects of decision measurement that need to be incorporated into the model: absolute and relative soundness.

Figure 1 shows schematically the metric $M(d,s(t))$ of decision soundness. M is a function from $D \otimes S$ to the real numbers that provides the figure of merit for the decision taken given either the actual (absolute) or perceived (relative) situation. The details of the definition of M must be tailored to match the purposes of each individual study. However, in general terms, M can be conceived of in terms of 'distances' between two situations, which we represent with the notation $\|s_1, s_2\|$. If a 'distance' metric can be defined on S, then the two types of soundness measures can be represented as

Deleted: Figure 1
Inserted: Figure 1
Deleted: Figure 1

$$M_{\text{abs}}(d, s_c(t)) = \left\| s_c(t) + \delta(d, s_c(t)), S_g \right\|, \tag{4}$$

based on true situation knowledge, and

$$M_{\text{rel}}(d, \hat{s}(t)) = \left\| \hat{s}(t) + \delta(d, s_c(t)), S_g \right\| \tag{5}$$

based on the decision-maker's situation knowledge.

The distinction between M_{abs} and M_{rel} concerns the decision-maker's situation awareness. M_{rel} is calculated by adding the implementation function δ to the perceived situation $\hat{s}$ rather than the actual situation s_c. This gives the quantity by which the decision-maker judges the soundness of decision options. On the other hand, M_{abs} gives the true soundness in terms of the actual outcome of implementing the decision.

Possibilities for defining the 'distance' between two situations include the following:

- The metric might be a quantification of the physical distance between the two situations, but this ignores the temporal, cognitive and C^2 dimensions of 'situation', such as the size of δ required to change s_1 into s_2.
- 'Distance' may be taken as the time required to change from one situation to another. This has the advantage of representing the time required to modify commander's intent, or to acquire situational awareness, but the metric might be unhelpfully skewed by physical-space issues, such as the time needed to make a transit.
- At a more abstract level, it may be useful to take the distance metric as the inverse of the probability that a path between the two situations exists and will be chosen. That is, if one can in some sensible way enumerate all possible paths from the current situation s_c, one can ask what fraction of them lead to the goal set S_g, and use the reciprocal of that quantity as the metric. For example, if S_g includes the requirement that a certain platform survive that has just been heavily damaged, then the probability of a path through the situation space to S_g becomes very low, and indeed the situation space may have become disconnected.

Performance Modulators and Biases

Effects of performance modulators and cognitive biases can be represented as perturbations of the entities defined above. For example:

1) a shift in $F_{CI}(s(t),R)$—i.e., in the decision-making function

2) a shift in $\hat{s}(t)$—i.e., in the understanding of what the situation is

3) a shift in $S_{CI}(t)$—i.e., in the commander's intent.

4) a shift in R—i.e., in the promulgated ROEs or other constraints.

If $d_{CI} = F_{CI}(s_c(t),R)$ represents the best decision relative to the extant commander's intent, then the actual decision d taken under the influence of one or more of the four shifts above can be written as $d = d_{CI} + \gamma$, where γ is a random fluctuation representing a difference between two decisions. Whether γ has a mean of zero depends on the effect being modelled. For example, many stressors are deleterious, so the mean of γ representing such effects would be non-zero. On the other hand, fluctuations intended to contrast analytical and recognition-primed decisions would probably have a mean value of zero and different variances. Similarly, a bias may cause a decision-maker to misunderstand what the current situation is, leading to a perturbation in $s = s_0 + \delta$. A shift in the understanding of the commander's intent could represent a misunderstanding of whether $s_c \in S_{CI}$ or in the application of the constraint set R. If the results of situational assessment ($\hat{s}$) and decision-making are modelled by probability distributions, then biases may be specified by shifts in the mean or variance of the distribution around the true value. A shift in the understanding of the scope of the commander's intent might also be modelled as a barrier function, such that the probability of a situation being labelled as in or out of the set follows a specified distribution. A bias might then be to interpret some aspects of a situation more tightly or loosely (more or less likely) in terms of whether it meets the criteria of the commander's intent. A bias may cause a similar interpretation on how strictly constraints are applied in the decision function.

Implementing the Conceptual Model

Key Implementation Issues

The conceptual model detailed above is not meant to be definitive, but to provide an alternative means to describe and think about decision-making modelling. To move beyond concept to implementation requires the following:

1) defining the attributes of the elements of the situation space S
2) defining the attributes of the elements of the decision space D
3) determining $F_{CI}(s_c(t),R)$—decision option generation and choice
4) determining $\delta(d,s_c(t))$—how decision d changes situation $s_c(t)$
5) specification of the impact of biases on each of these elements
6) determination of the form of M, the metric for decision quality.

Development of the Model

To start implementing the decision-making model we make the following assumptions:

1) Elements of the situation space can be classified and expressed as an array of characteristics.
2) The decision space is enumerable and expressible as an array of characteristics.
3) Decision-option generation is possible using an evolutionary algorithm with mutation, incorporating decision-maker understanding of commander's intent and decision-maker biases.
4) Decision choice can be based upon battlespace simulation which incorporates decision-maker beliefs and biases.
5) Situation change as a function of a decision implementation can be based upon the true situation and true unit capabilities, and evaluated using a single step of a battlespace simulation.
6) A relative decision soundness metric can be defined using the output of a battlespace simulation to calculate the increase or decrease in likelihood of achieving commander's intent.

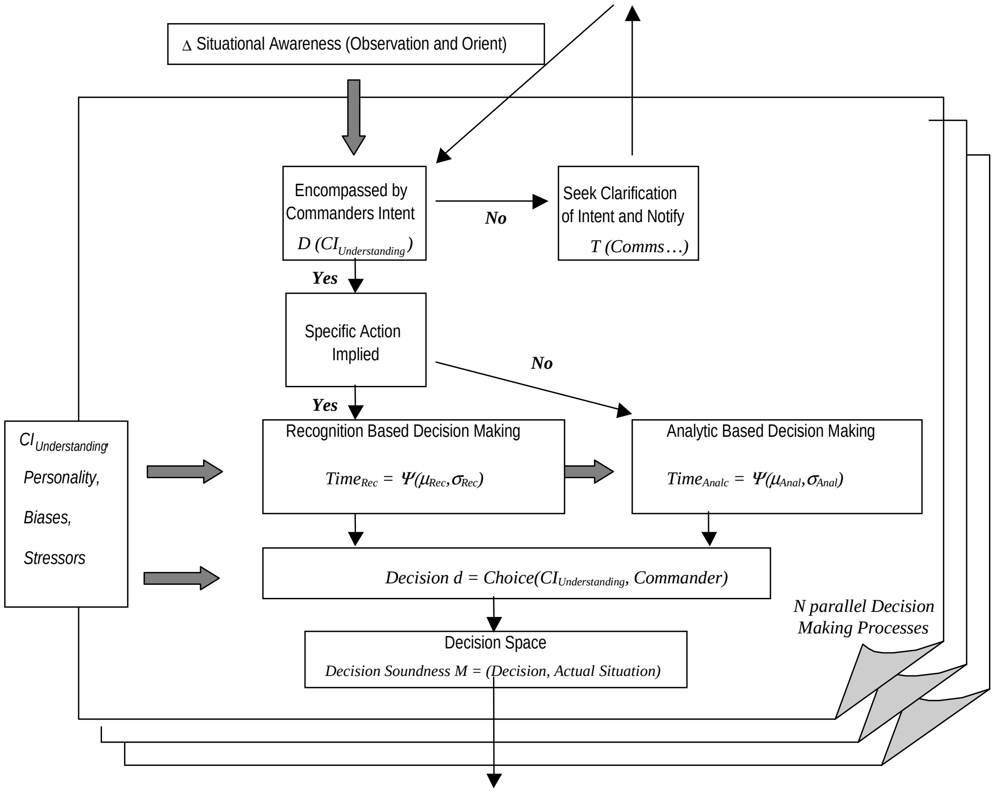

Figure 2: Decision-making model components and flows. Boxes are individual processes, with specific calculations given in italics.

Based on these assumptions, Figure 2 shows a more developed view of the modelling framework, and one from a different perspective than Figure 1. It explicitly includes a set of independent and parallel decisions, and a mechanism for altering the commander's intent in the event of a situation being encountered that is believed to lie outside of the extant intent (outside S_{CI}, in terms of Figure 1). The steps through the model are conceived of as follows:

Deleted: Figure 1
Inserted: Figure 1
Deleted: Figure 1
Deleted: Figure 1
Inserted: Figure 1
Deleted: Figure 1

- The decision-maker receives some change Δ_{SA} in situational awareness.
- The Δ_{SA} is compared to the commander's intent (CI), as understood by the decision-maker. If the Δ_{SA} is covered by the CI then the decision-maker moves down the process; if the Δ_{SA} is not covered, then the decision-maker must seek clarification and also notify higher command. This will incur a cost in time (T(Comms…)). In some cases, this part of the process might actually require a number of iterations before a new CI is obtained that allows the decision-maker to progress.
- Once the Δ_{SA} is covered by the CI, the decision-maker then checks to see if the CI implies a specific action or decision (i.e., recognition-primed decisions); if not, an analytical decision is required. (The time for these events is determined in the next step.) This process will be a function of understanding of the CI, experience, personality, biases and stressors and may include a stochastic element.

- The means and variances of the distributions for decision speed and absolute decision soundness are determined using as inputs CI, understanding, decision-maker's experience, personality, biases and stressors, and whether the decision is analytical or recognition-primed.
- The calculated means and variances are then used to determine a decision point within the decision space and the decision speed.
- The soundness value, which is calculated given the decision point, and the decision speed together constitute decision quality.

The model clearly requires a quantification of the extant commander's intent and the decision-maker's situation awareness, if not separately, then at least so far that the degree of correlation between them can be represented by a number. At an abstract level, the degree of correlation could be taken as a random variate; alternatively, it may be determined from more detailed modelling.

Option Generation using Genetic Algorithms

The important feature of the decision-making model for present purposes is to capture how the decision-maker navigates the decision space and the process by which a decision-maker chooses a particular option, rather than the exploration of novel solutions from uncharted areas of the decision space. From this perhaps subtle point, it follows that the modeller can assume that all possible decisions are expressible and can be evaluated for soundness within the model. The same point can be viewed in another way, in terms of controlling for confounding effects: since the aim is to explore effects of network centricity, decision-maker creativity must be held constant except where influenced by a network-based application such as collaborative planning tools.

This raises the question of modelling decision-option generation, at first sight the least tractable of the modelling issues. There have been a number of studies using evolutionary methods such as genetic algorithms to generate search plans and other tactical decisions. [12] Given that the decision options can be quantified to the point that they can be represented by a set of binary variables, a genetic algorithm can be applied to develop 'optimal' sets of decisions. By 'optimal', we do not mean necessarily globally optimal, since proving the convergence of genetic algorithms to a global optimum requires the imposition of undesirable constraints on the algorithms. Instead, we contend that the population of solutions developed models well the result of the satisficing process that is a significant and widespread feature of military decision-making [13]. Satisficing is the process of looking not for an optimal solution but for one sufficiently good enough to work. Indeed, the entire process of evolutionary algorithms is an attractive computational metaphor for human option generation in military decision-making [14]. As we use the terms, the 'evolutionary' or 'genetic' aspects of the process include the following characteristics:

- Fitness/Reproduction – some candidate options become more favoured in the population (even if they do not satisfice the situation, either because of a bias toward them or because there is a part of them that has a potential and should not be lost).
- Cross-over – a candidate option is modified by the incorporation of part of another solution (mixing and matching good looking parts of solutions).
- Mutation – the random incorporation of new ideas into an option, and the random drift and

variance of ideas over time.

As applied to modelling option generation, a genetic algorithm would have the following steps:

1) An initial population of decisions is developed based upon the decision-maker's experience, biases and understanding of the problem constraints.
2) The decision population is evaluated against a set of criteria for soundness.
3) If a solution that satisfices the criteria is not found, then new solutions are developed by
    a. Combining parts of good solutions from the old population
    b. Mutating some of the attributes of the old population.
4) The new population is evaluated for soundness and the process is iterated until a satisficing solution is encountered.

To take this further requires that the commander's biases and experiences be quantified as well, so as to match particular situations with potential decision options that have an associated likelihood of being chosen by a particular decision-maker. Similarly, decision-maker biases toward certain types of decisions can be included both in the selection of an initial population and as a weighting on the perceived decision soundness and ability to satisfice the situation.

Decision Soundness and Quality

The decision soundness should be a measure of the likelihood that the decision will lead to the attainment of the operation's goal, that is, the commander's intent. A decision that increases the likelihood should have a higher soundness than one that decreases the likelihood. Additionally, this likelihood can be relative to the likelihood under the decision-maker's best knowledge, or relative to the true likelihood.

Evaluation of decision soundness, both to satisfy the decision-option-generation process and as part of the overall decision quality metric, is a key part of this model. This is a difficult task; in essence it requires evaluating the effect of a decision on the future probability of achieving the commander's intent. In the case of decision-maker choice, it should include decision-maker biases and beliefs; as part of the absolute quality metric, the evaluation must be free of those biases and reflect the true likelihood. In either case, even if a simulation may be structured to estimate the outcome of a decision made at time t_i on the situation at t_{i+1}, the resulting uncertainty of actions at t_{i+1} will result in an exponentially growing tree of possible situations over time. A complete evaluation of likelihood will require modelling the outcome of all the sets of decisions—a huge task.

In terms of modeling the decision-maker's assessment of the option satisficing, it is probably enough to make a rough estimate of the likelihood based upon a small number of model runs. Increasing the number of runs to provide a Monte Carlo-esque estimate of likelihood with increasing experience might even simulate commander experience. Likelihood estimates can be further estimated utilizing synthetic annealing methods to sample the decision space [16]. For the purposes of the task at hand, it appears clear that only simplistic situations are likely to be completely tractable. Alternatively, a lower-fidelity but more straightforwardly computable situation-space metric of the element-by-element difference between the new situation and the goal of the commander's intent might be sufficient in some cases.

As numerous remarks throughout this paper make clear, we have no illusions about the difficulty of implementing the program laid out above. Each of the key implementation issues is likely to be problematic to practical implementation. The use of software agents coupled with evolutionary decision-option generation for social and economic modeling [14,15] provide some indication that it is possible.

Example: Applying the Conceptual Model to Maritime Interdiction Operations

In [1], an attempt to quantify the value added of NCW to Maritime Interdiction Operations (MIO) was reported. A major gap in that analysis was the inability of the queueing theory military effectiveness model to definitively link the addition of any particular network application to increases in military capability. All links made were of a qualitative nature. While the decision-making model described above has not been developed to the point of quantitative results it is instructive to apply the concepts to the particular warfare problem.

As an example, the MIO analysis examined the problem of an interceptor force defending a coast-line from smugglers. In the scenario analyzed, a finite number of interceptors are each given a section of coastline to defend and must stop and interrogate all shipping to and from the coast. The smugglers are expected to use a strategy of rushing the line of interceptors in order to overwhelm the units in a particular sector. The question posed was, if the commander had prior knowledge of where the rush would take place, perhaps from intelligence or some form of reachback, then could the force be reconfigured to handle the problem. Conceptually, it was assumed that, under NCW, when the commander obtained prior warning of the increase in smugglers (or traffic) in an area, it would be possible to use net-based collaborative planning to pull enough interceptors into the relevant sector to handle the increased traffic. Queueing showed that under these assumptions, the force could indeed maximise its probability of catching the smugglers.

In fact, the queueing theory actually shows that it is equivalent if the workload of the ships dealing with the rush can be reduced to a baseline level, and therefore the more fundamental decision for the force commander is the criteria to be used in determining how many vessels must be intercepted, and how to position sufficient interceptors to do it.

To apply the conceptual model the six key implementation issues must be addressed:

1) The situation space must be enumerated to account for the disposition and state of the interceptors, neutrals and smuggler units, and the MIO commander's understanding of the status of each. The physical state of the units will impose some conditions on the timing of information requirements. It is reasonable to assume that the MIO commander will have good situational awareness of the state of his own units, so that $\hat{s} \approx s_c$ so far as own forces are concerned. The question of whether a given other vessel is neutral or smuggling is the key information management issue in this scenario and is the most likely source of divergence of $\hat{s}$ from s_c. The situation space should be explicitly constructed to facilitate the depiction and exploration of this divergence.
2) The decision space must be determined, including the time scale and scope of the decision options. Since the emphasis is on modelling the effects of NCW, decision options must reflect the effect of various amounts of information on the intentions of incoming vessels, the manner in which the information is generated and passed between the entities of the model,

and the use that the entities make of it [2].

3) The option generation process must be seeded with the current operations plans, doctrine and tactics to provide an initial population of options. The commanders' intent and resources available are to be instantiated as constraints on the problem, possibly the current disposition of interceptors and other physical limitation as well. Recent modelling of MIO using a prioritised queue [17] indicates that the key decision issue is that of 'pre-emption'—the conditions under which the MIO commander decides to abandon the interception of a low-priority vessel in order to free up an asset for intercepting a recently arrived higher-priority vessel. It is important that the modelled decision space include a representation of pre-emption and desirable that several pre-emptive options be generated wherever possible and appropriate.
4) Physical limitations on interceptor ability to respond and implement a decision option, and likely smuggler responses if a multi-step (response-counter) study is to be attempted, should be included in the definition of the implementation function δ. However, to match the previous MIO study, a single-step response is sufficient.
5) To simplify the problem, we leave out the issue of biases, although the issue of complacency for units who have been on patrol for long periods might be addressed by including a bias in the option generation functions.
6) The measure of effectiveness is the probability of intercepting a smuggler. This will be determined through a combination of the number of interceptors made available and their effectiveness as given by the queueing theory model.

The application of the conceptual model thus focuses the problem on the type of information required by the decision-maker and therefore on specific NCW applications that can affect the information. On implementation, specific issues of information quality (clarity, timeliness etc.) can be explored.

Summary and Conclusions

A framework for a model of decision-making is developed. This involves the construction of two abstract spaces, one representing all situations and the second all decisions. Functions connecting these are constructed: the function from situations to decisions describes the decision-making process; that from decisions to situations describes the implementation of the chosen decision option. This process has the cyclical nature of the OODA loop. The two spaces also support a function quantifying the soundness of the chosen option. Several general suggestions are presented for suitable soundness metrics; a particular metric must be chosen to suit the problem at hand. The conceptual framework is developed sufficiently to point the direction for implementation, to give an indication of expected outputs and capabilities, and to suggest its relevance for the modelling of NCW.

The model incorporates most of the important influences on decision-making from the literature and provides at the very least a conceptual model within which to think about these issues. Despite the formidable obstacles to full implementation, we believe that the framework set out above is sufficiently detailed to provide an indication of the sort of outputs that are achievable. The most important of these is the ability to assess the impact of the inputs to decision-making on force effectiveness as embodied in the commander's intent. This would provide a basis for judging the effect of varying levels of network centricity.

Further, even if a numerical implementation should turn out to be impractical, the conceptual model of provides a formalism within which the functional relationships of changes to decision-making inputs can be constructed and studied. Thus, to take the effects of biases as an example, in cases where it proves impossible to quantify the effect, it may be possible to see where a bias influences the process. This would then assist users to construct human-in-the-loop experimentation designed to explore the specific issue at hand.

References

[1] M.G. Hazen, R. Burton, R. Klingbeil, K. Sullivan, M. Fewell, I. Grivell, C. Phips and P. Marland *Modelling the Effects of NetCentric Maritime Warfare (NCMW) in Maritime Interdiction Operations (MIO).* 8th International Command and Control Research and Technology Symposium, National Defence University, Washington DC, 17-19 June 2003

[2] M. Fewell and M.G. Hazen *Network-Centric Warfare: Its Nature and Modelling*, Defence Science and Technology Organisation, Australia. DSTO Research Report 0262, September 2003

[3] R. Pigeau & C. McCann (2000) 'Redefining command and control' in C. McCann & R. Pigeau (eds) *The human in command: exploring the modern military experience*, NY, Kluwer Academic, pp. 163–84.

[4] M. Fewell and M.G. Hazen (2004) *Cognitive Issues in Modelling Network-Centric Command and Control*. Defence Science and Technology Organisation, Australia. DSTO Research Report (in Review)

[5] R. Klingbeil (2001) 'An analysis of command decision time delay' in *Proc. 1st meeting of MAR Group AG-1*, TTCP document DOC-MAR-1-2001, vol. 2.

[6] C.D. Wickens & J.G. Hollands (2000) *Engineering psychology and human performance*, 3rd edn, Upper Saddle River NJ, Prentice Hall.

[7] Von Neumann & O. Morgenstern (1953) *Theory of games and economic behavior*, 3rd edn, Princeton NJ, Princeton University Press.

[8] M.R. Lehto (1997) 'Decision Making' in G. Salvendry *Handbook of human factors and ergonomics*, 2nd edn, New York, John Wiley & Sons.

[9] D.W. Dorsey & M.D. Coovert (2003) 'Mathematical modeling of decision making: a soft and fuzzy approach to capturing hard decisions', *Hum. Fact.* **45**, 117–35.

[10] D. O'Hare, M. Wiggins, A. Williams & W. Wong (1998) 'Cognitive task analyses for decision centred design and training', *Ergon.* **41**, 1698–718.

[11] M.R. Endsley (1995) 'Toward a theory of situational awareness in dynamic systems', *Hum. Fact.* **37**, 32–64.

[12] D.P. Kierstead & D.R.DelBalzo (2003) 'A Genetic Algorithm Approach for Planning Search Paths in Complicated Environments', *Mil. Op. Res.* **8**(3), 45.

[13] M.S. Cohen, J.T. Freeman & S. Wolf (1996) 'Metarecognition in time-stressed decision making: recognising, critiquing, and correcting', *Hum. Fact.* **38**, 296–19.

[14] E. Chattoe (1998) 'Just How (Un)realistic are Evolutionary Algorithms as Representations of Social Processes'. *J. Art. Soc and Soc Sim* **1**(3) [www.soc.surrey.ac.uk/JASSS/1/3/2.html]

[15] J. Rees and G. Koehler (2002) 'Brainstorming, Negotiating and Learning in Decision Group Support Systems: An Evolutionary Approach'. *Journal on Computing*, **14**(3), 278-292.

[16] C.A. Gillard, D.J. Thomson, and G.J. Heard (2003) 'Estimating geoacoustic parameters using matched-field inversion methods', *IEEE J. Ocean Eng*, **28**(3), 394- 413.
[17] I. Grivell and M.P. Fewell (2003) 'Towards an End-to-End Model of Maritime Interception Operations' presented to the Portsdown meeting of Action Group 1 of the Maritime Systems Group of The Technical Cooperation Program (October, 2003).