

**ARMADA DEL ECUADOR
ACADEMIA DE GUERRA NAVAL
Guayaquil**

-0-



LECTURAS RECOMENDADAS

“Predictive Analytics in the Naval Maritime Domain”

Bonnie Johnson

Naval Postgraduate School Systems Engineering Department

Lectura Recomendada por:

CPCB-EM Othoniel Palacios Celín
Oficial Instructor de la Academia de
Guerra Naval

2023

Predictive Analytics in the Naval Maritime Domain

Bonnie Johnson

Naval Postgraduate School Systems Engineering Department
bwjohnson@nps.edu

Abstract

Predictive analytics offers a potential game changing capability for Naval tactical decision superiority. Tactical operations could take a significant leap in progress with the aid of a real-time automated predictive analytics capability that provides predictions of second and third order effects of possible courses of action. This future capability would accompany current developments in the use of artificial intelligence and data analytics to improve battlespace knowledge and offer automated battle management aids to the tactical warfighter. As the automated battle management aids develop tactical course of action options the predictive analytics capability could predict how the adversary might respond to each course of action option. The predictive analytics capability could continue to “wargame” possible blue force/red force actions and responses—generating predictions of second and third order effects. These predictions offer the tactical warfighter a more strategic perspective in making tactical course of action decisions. By performing this analysis using an automated aid with artificial intelligence, it allows the capability to support real-time decisions and to analyze great amounts of data (both sensor data and historical data) and handle highly complex tactical environments. This real-time wargaming translates into high order computations that would be impossible to be performed manually in the short reaction times given. This paper discusses the results of a study of predictive analytic capabilities in the naval maritime domain.

I. Introduction

A predictive analytics (PA) capability – that can take into account possible consequences and effects into the process of decision-making – is key to enabling decision superiority for naval forces. The PA capability, based on automated data analytics, would support battle management aids (BMA) by developing “what-if” and “if-then” predictive scenarios to shape the synthesis of future intelligent decisions and adaptive capabilities. This conceptual capability would inform decisions concerning courses of action (COA) based on what the longer-term effects are projected to be. It would enable short-term and long-term objectives to be weighed as tactical decisions are made. In essence, a PA capability is a critical step in enabling a real-time wargaming capability for naval operations.

The goal of decision making is to change the probabilities of outcomes to make preferred outcomes more likely. The

decision-maker’s choices act upon the world, causing changes in outcome probabilities (Cox 2015). Yet, this crucial concept of causal efficacy is seldom developed in detail in decision analysis, and the fact that formal probability theory applies only to events rather than to actions and their consequences is seldom emphasized (Pearl 2008). The use of Bayesian Networks (BN) and causal graph models may provide a solution to predict probabilities of outputs given inputs and observations. These types of models can be used to build quantitative representations of complex dynamic situations. Dynamic BN models and BN-learning algorithms can learn from data to create an adaptive capability that can predict outcomes in a changing environment.

Using methods of machine learning to process and analyze large amounts of heterogeneous data and information, artificial intelligence (AI) technology can make predictions about probable effects, outcomes, and responses. These PA and AI methods can provide a powerful capability for tactical decision-making. Armed with the knowledge of possible effects and adversary responses to courses of action, warfighters can leap ahead in terms of applying longer-term strategy to near-term warfare decisions. A critical enabler of developing an executable model of blue forces and red forces is the incorporation of the correct metrics, premises, and assumptions (Talbot and Ellis 2015).

This paper begins (in Section II) with a description of the authors’ concept for a future predictive analytics capability that could support a real-time operational automated decision aid. Section III discusses data concepts required to support such a future PA capability. Section IV contains an overview of AI and game theoretic methods that show promise for enabling an automated PA decision aid. Finally, Section V contains the conclusion.

II. A Conceptual Naval Maritime Predictive Analytics Capability

The ability to perform predictive analytics in support of maritime operations, such as planning and tactical warfare, requires a set of analytic capabilities that study the available data, develop COA options, and make predictions concern-

ing their effects for the purpose of selecting options with desired effects. Figure 1 illustrates a conceptual framework for a PA capability for the naval maritime domain.

Required inputs to this capability are shown as self-awareness and situational awareness knowledge. Self-awareness amounts to the development of a blue force model which keeps track of the location, status, and capabilities of the blue force resources or warfighting assets. Situa-

Each of the 1st order effects is then analyzed (step 4), based again on the red force model, to estimate a set of possible adversarial responses. These constitute 2nd order effects. Each 1st order effect may map into one or more possible 2nd order effects.

The 2nd order effects, which may now contain a significant number of possibilities, are analyzed (step 5) using knowledge of our blue forces (contained in the blue force

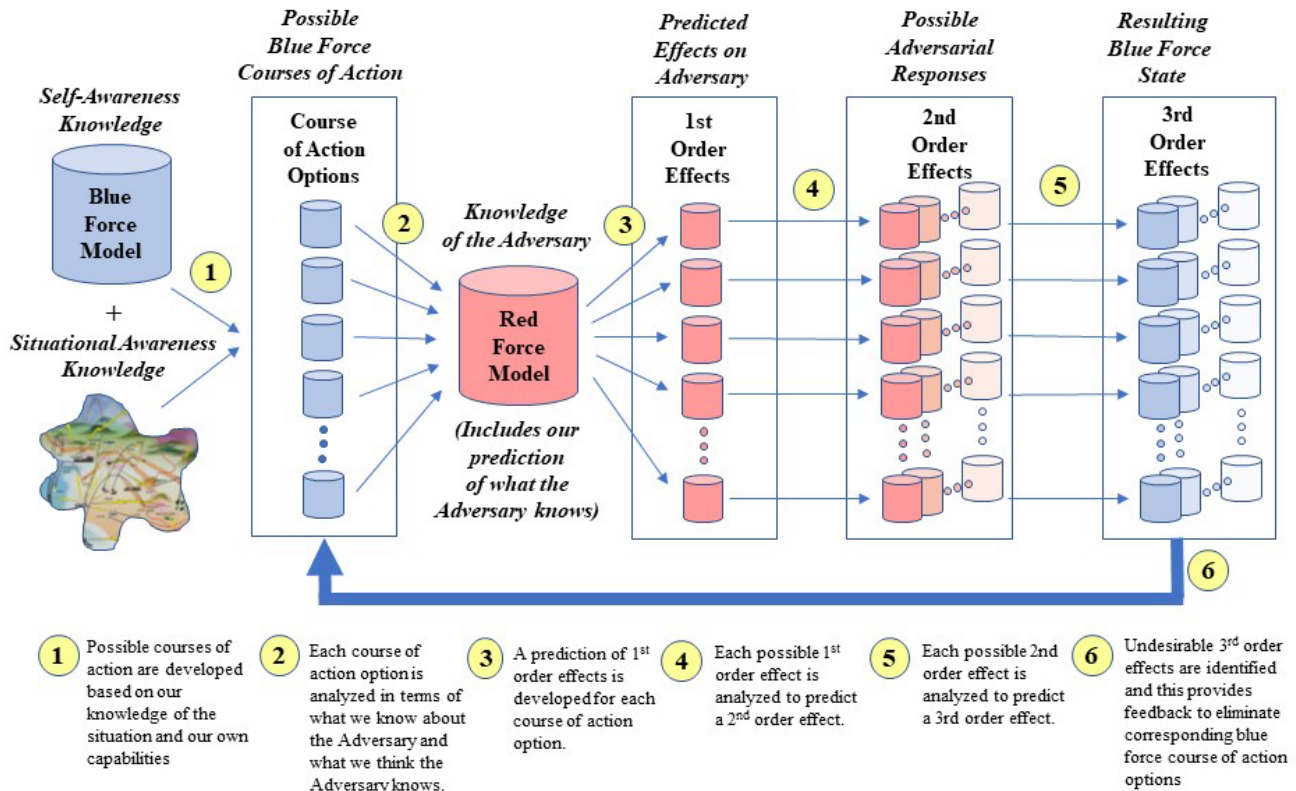


Figure 1 – Conceptual Framework for Predictive Analytics Capability

tional awareness consists of real-time sensor data feeds that are fused and analyzed to provide an understanding of the battlespace or operational environment. From the blue force and situational awareness models, a set of possible COAs (shown as step one) are developed that represent a continuum of possible blue force actions that can be taken at any moment in time. These include, as examples, the placement or movement of assets, sensor tasking, weapon engagement decisions, the and the use of countermeasures.

The capability requires a model of the red force, or adversary, that estimates what is known about the adversary as well as predicts what the adversary knows about the blue force and the situation. The PA capability evaluates (step 2) each COA option in terms of our knowledge of the red forces to predict (step 3) the effects of each option on the adversary. Steps 2 and 3 produce a set of predicted 1st order effects. Each blue force COA option has a direct mapping to its predicted 1st order red force effect.

model) to predict the 3rd order effects. The 3rd order effects are a set of predicted blue force states that result from the possible adversary responsive actions. Thus, there is a one-to-one mapping of possible 2nd order effects to possible 3rd order effects.

The set of 3rd order effects are evaluated (step 6) to identify undesirable outcomes. Any undesired 3rd order effects can be used to feedback into the set of blue force COA options and eliminate undesired COAs. Thus, the conceptual PA capability is an analysis tool to provide a deeper understanding of the COA options in terms of their possible causal effects and expected consequences.

Each step in the PA capability can include an estimate of the certainty of the analysis, providing a level of confidence in the predictions. This would add even greater refinement in terms of evaluating the desirability or undesirability of 3rd order effects, and consequently blue force COA options.

The conceptual PA capability can enhance future automated tactical decision aids. Figure 2 illustrates a tactical decision aid, showing how capabilities for PA and knowledge discovery would interact with the tactical resources (shown along the bottom row) as well as the “decision engine.” The conceptual resource management capability would assess and prioritize missions and use those results to develop the COAs, which the PA capability would evaluate based on predicted 1st, 2nd, and 3rd order effects.




Figure 2 – Predictive Analytics as a Core Capability for an Automated Tactical Decision Aid

III. Data and Knowledge Concepts for Predictive Analytics

A naval tactical “decision-maker is not interested in data or big data as such; but the knowledge it provides. (Zhao, Kendall, and Young, 2015, p. 22).” They are interested in actionable knowledge required to gain and maintain the tactical advantage. Gaining and maintaining knowledge of the maritime domain is not only a required capability that enables the conceptual PA capability, but it has a direct impact on the accuracy of the predictions made. The levels of completeness and accuracy dictate how good the internal models are as well as the predicted 1st, 2nd, and 3rd order effects.

Three categories of tactical maritime knowledge are illustrated in Figure 3 as knowledge of the blue forces, knowledge of the red forces, and knowledge of the operational situation. Conceptually, computer-aided models of each could be created to support real-time naval operations as well as the PA capability. Each model could be developed and continuously updated based on the input data that is constantly changing to reflect the changes in the states of the blue forces, the red forces, and the environment.

The blue force model would represent all that is known about the blue forces at any given time. It would provide the Navy with self-awareness by containing what is known

about the status, location, and readiness of the blue force warfighting resources or assets. The model would assess the readiness of each resource as well as the overall force readiness. The model would contain an assessment or prediction of each resource’s capability to perform an assigned COA. Examples include probability of kill, probability of detection, probability of jamming, etc. The model could also predict overall force capability given a particular threat environment.

The red force model is envisioned as an estimated prediction of what is known about the adversary based on data and intelligence available. This model would estimate what types of capabilities the red force possesses and approximate the overall red force readiness. The model would predict the adversary’s intent, tactics, and strategies for the purpose of predicting how the adversary might act in different situations or respond to blue force actions. The model could make an educated guess as to what the adversary knows about the situation and about the blue forces. This prediction would be based on an assessment of the blue force’s possible visibility to the red force based on what is known about the adversary’s location and surveillance capabilities. The red force model would become the Navy’s prediction of what is known about the adversary.

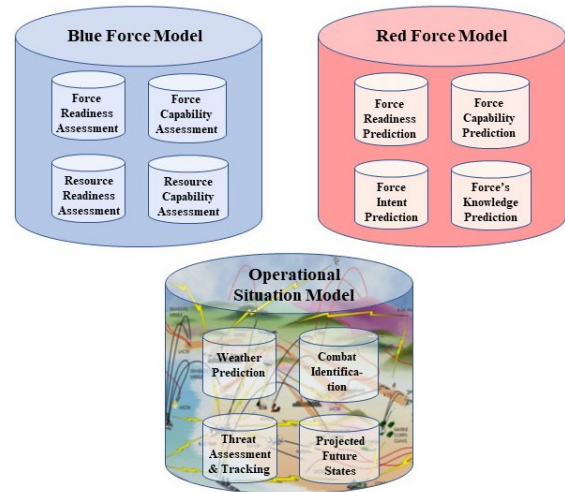


Figure 3 – Conceptual Computer-Aided Models as Enablers for a Predictive Analytics Capability

The operational situation model would constitute the current situational or maritime awareness. This model, based primarily on real-time sensor data, would contain the understanding of the battlespace in terms of the weather, combat identification, and threat assessment and tracking. It would be comprised of information on the location, kinematics, and identification of all objects (friendly, neutral or foe) in the area of interest. The model’s completeness, accuracy, and up-to-date-ness would depend entirely on the data collected. The operational situation model would also contain

predictions of projected future states of the area of interest. Examples of this could include projected impact points of threat missiles, projected locations of enemy aircraft and ships, and future weather and environmental conditions.

Developing and maintaining these changing and actionable models depends on a number of data collection, fusion, security, and management capabilities. The naval tactical domain has data architectures in place for collecting, processing, and fusing sensor data for developing situational awareness of the battlespace. This data supports combat identification, threat identification and tracking, and projections of kinematic objects in the area of interest.

In order to develop an internal model of the blue force, the Navy would also need to collect data concerning blue force asset status, location, and capability (Johnson, 2019, Rowe 2019). Brown (2019) developed a conceptual architecture for collecting blue force asset data to support the determination of force readiness. Self-awareness data could also be used to determine individual resource readiness and general blue force self-awareness.

In order to develop an internal model of the red force, the Navy needs to analyze the situational awareness data along with information from intelligence sources to make inferences about the capabilities, location, and readiness of red force assets. The use of intelligence sources could be used to model likely red force intent, tactics, and strategies. The combination of predicted red force asset knowledge with knowledge of our blue force assets can be used to make inferences about the adversary's knowledge of the blue force.

Maintaining knowledge of the real-world is a critical part of implementing a PA capability for the tactical Navy. The models provide a "belief state" that become the basis for making predictions about the consequences of COAs. Russell and Norvig (2010, p. 480) write that the belief state is a "representation of the set of all possible worlds" that a system may exist in. The belief state is then used to generate COA options and corresponding possible outcomes and consequences, and to evaluate these options.

IV. Artificial Intelligence and Game Theoretic Methods for Predictive Analytics

A number of data analytic methods exist that can support the many different types of estimation and predictive capabilities that have been described up until this point. For example, Kalman filters have been widely used for projecting the future kinematic states of moving objects in the battlespace. This is a form of computational prediction. Data fusion analytics are used to combine and assess heterogeneous data from different types of sensor to enhance our ability to identify and understand combat objects in the battlespace. This section focuses on AI and game theoretic methods that can potentially be used to evaluate the COA options by predicting 1st, 2nd, and 3rd order effects.

A. Predictive Analytics as a Data-Driven and Automated Process

Abbott (2014) describes predictive analytics as a data-driven process of discovering interesting and meaningful patterns and inducing models from the data, rather than basing results on assumptions made by the analyst. PA is a process that results in discovering variables to be included in the model, parameters that define the model, weights or coefficients in the model, and also the very form of the model itself. These models can then be used to build predictions. PA, as described by Abbott (2014), does not do anything that a human analyst could not accomplish manually given enough time. The reason to automate the process is because the number of variables and permutations can quickly result in thousands of computations. Automated algorithms can sift through the many potential combinations of data to identify patterns and interesting results.

One aspect of the conceptual PA capability that is beyond human capability is the ability to store or "remember" all of the numerous COA options, permutations, effects, and outcomes. In a tactical combat situation, these options and effects would be changing continuously as the environment changes, creating an even more complex memory challenge. And for the envisioned PA capability, the 1st, 2nd, and 3rd order effects and permutations needs to be stored and easily accessed for evaluation.

Automation also plays an important role in providing data fusion, correlation and analytics for developing and continually updating the blue force, red force, and operational situation models required for the conceptual PA capability.

B. Statistical Methods

Statistical methods are widely used to perform a confirmatory analysis concerning a hypothesis involving a relationship between inputs and outputs (Abbott 2014). The analysis confirms or denies the causal relationship and quantifies the degree of that confirmation or denial.

Regression analysis is a statistical method for analyzing and modeling the relationship between a continuous dependent variable and an independent variable to build a model for making predictions. The first step is to identify and explain the best model that represents the relationship between the dependent and independent variables. The second step is to use this model to predict future values of the dependent variable given specific values of the independent variable. (Kalaian and Kasim 2017)

Discriminant analysis is a statistical technique that uses the information from a set of independent variables to predict the value of a discrete (categorical) dependent variable, which represents the mutually exclusive groups in the predictive model (Kalaian and Kasim 2017). Discriminant analysis can be used to identify the best combination of independent variables or predictors, that provide the best discrimination between groups in an effort to accurately predict

a membership in a particular group. This technique can be used for threat identification (as friendly, neutral, or foe) to match a tracked object’s characteristics to the appropriate group’s predictive model.

Table 1 lists differences between using statistical methods and PA methods. Statistical methods can apply to small data sets and rely heavily on ensuring the models are built properly and are typically linear; whereas PA methods draw heavily on machine learning and AI, require lots of data, and have no provable optimum solution.

Table 1 – Statistics vs. Predictive Analytics (Source: Abbott 2014)

Statistics	Predictive Analytics
Models based on theory; there is an optimum.	Models often based on non-parametric algorithms; no guaranteed optimum.
Models typically linear.	Models typically nonlinear.
Data typically smaller; algorithms often geared toward accuracy with small data.	Scales to big data; algorithms not as efficient or stable for small data.
The model is king.	Data is king.

C. Graph Theory Methods

Bayesian networks (Bayes network, belief network, decision network, Bayes model, or probabilistic directed acyclic graphical model) are a category of statistical models that represent a set of variables and their conditional dependencies in the form of graphs. Bayesian networks offer a systematic way to represent relationships between variables and their dependencies explicitly and concisely—greatly simplifying the process of specifying probabilities for the large numbers of variables that may exist (Russell and Norvig 2010).

Bayesian networks are ideal for taking an event that occurred and predicting the likelihood that any one of several possible known causes was the contributing factor. For example, a Bayesian network could represent the probabilistic relationships between diseases and symptoms. Given symptoms, the network could be used to compute the probabilities of the presence of various diseases. Efficient algorithms can perform inference and learning in Bayesian networks. Bayesian networks that model sequences of variables are called dynamic Bayesian networks. Generalizations of Bayesian networks that can represent and solve decision problems under uncertainty are called influence diagrams.

The network’s nodes represent observable quantities, hypotheses, or unknown parameters. The edges represent conditional dependencies. Each node is associated with a probability function whose input is a particular set of values representing the node’s parent variables and whose output is the probability of the variable represented by the node.

A Bayesian network representing the probabilistic variables and their relationships for the combat identification of an air object is shown in Figure 4. This network shows factors involved in determining the combat identification of an airborne object based on what information is known about the object and its environment. Factors, such as what is known about the object’s kinematics, the object’s proximity to the airport and the level of turbulence in the near environment, are shown as variables (or nodes) in the network. The network also contains nodes representing how the object is identified, such as by intelligence, interrogation friend or foe (IFF) or by electronic surveillance means (ESM). It can be noted that the directions of the arrows can be used to show the causal relationship between the actual identity (in the real world) and how it will affect the factors that allow it to be identified. The arrow directions can also be reversed (as in Figure 4) to show that given a variety of information sources, they can be used to support the identification of the object.

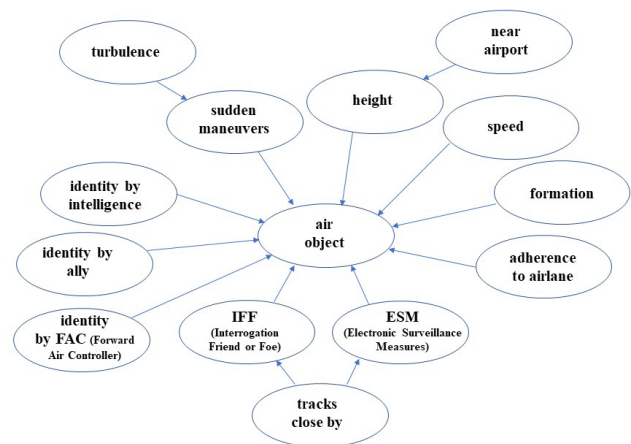


Figure 4 - An Example Bayesian Network for Combat Identification (Source: van Gosliga and Jansen, 2003)

D. Decision Theory

Decision theory provides methods for selecting among actions based on the desirability of outcomes, often in situations that are only partially understood. Russell and Norvig (2010, p. 610) describe these situations as “nondeterministic partially observable environments.” Thus, the AI system may not know the current state completely, so a random variable is used to represent the possible outcome states. The decision-maker’s preferences are represented by a utility function which assigns values corresponding to the desirability of the possible outcome states.

Automation can support the development and application of utility functions. The utility functions become complex for complex decision spaces, such as military tactical operations. Such problem spaces are characterized by many possible outcomes and many possible factors affecting each outcome. They also introduce uncertainty and dependences

among the variables representing factors. Automated systems can develop probability models reflecting the stochastic processes that generate outcomes. The systems must also model the error in the utility estimates that may be introduced by unknowns, incomplete knowledge, and bias. The use of multi-attribute theory along with the models of expected utilities and associated error can provide an automated aid for making decisions.

Decision theory can be thought of as the combination of probability theory and utility theory. The use of a decision network, also called an influence diagram, combines Bayesian networks with node types for actions and utilities. Decision networks provide a useful framework to aid AI in making complex decisions involving multi-attributes, multi-variables, many possible outcomes, and knowledge uncertainty. Figure 5 shows an example of an influence diagram with a military application. The oval nodes, referred to as chance nodes, represent random variables. The rectangle nodes, called decision nodes, represent decision points where there is a choice among actions. The hexagonal nodes are the utility nodes which represent the AI system's utility function.

E. Learning-Driven Methods

Learning in terms of AI systems is defined as “the capability of drawing intelligent decisions by self-adapting to the dynamics of the environment, taking into account the experience gained in past and present system states, and using long term benefit estimations” (Kim 2018, p. 222). Implementing learning algorithms requires large amounts of training data. Kim (2018) explains that progress is being made in learning algorithmic game theory which lies in the intersection of game theory and AI learning algorithms. These methods show potential for the military domain by implementing many iterations of a wargame and training the learning algorithms to identify the best COAs and blue force strategies based on desired game outputs.

Supervised learning (also referred to as predictive modeling) is a method that uses a “supervisor” target variable to represent the answer to a question of interest or a value that is unknown but could support decision-making if known. Supervised learning uses “ground truth” to train the AI system using prior knowledge. The goal is to learn a function given some input data and desired outputs that best approximates the relationship between the input and desired output.

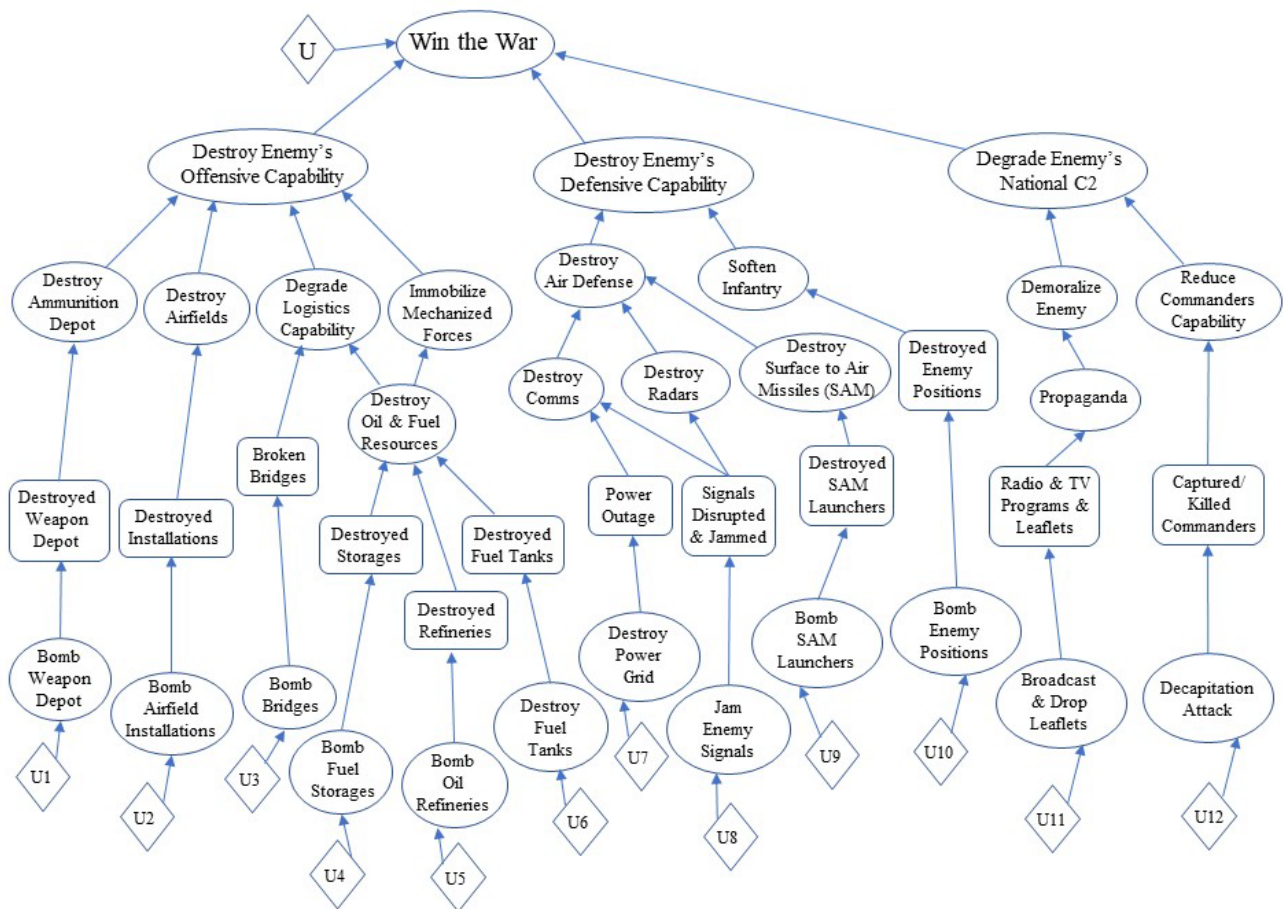


Figure 5 – An Example Decision Network for Military Actions (Adapted from Campos and Ji 2008)

This function can then be used to classify target variables or perform regression on continuous target variables.

AI machine learning methods can be used to support predictions based on comparing real-time data with “best models.” Figure 6 shows a process of first training the system to find a best model by running many iterations allowing supervised learning to occur. The best model can then be used in the operational system (in the second row) as a standard by which to compare incoming real-time data. As data begins to match the model, future state predictions can be inferred.

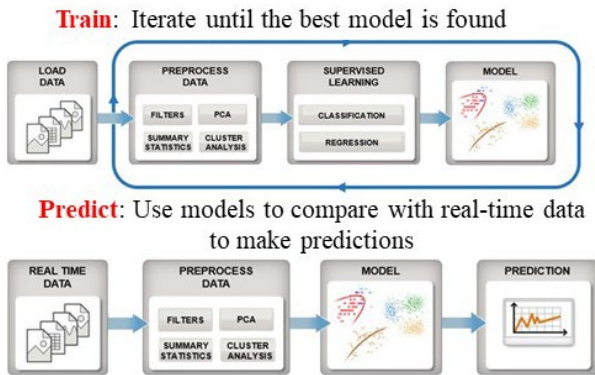


Figure 6 – Example of Supervised Learning for Predictive Analytics

Unsupervised learning (or descriptive modeling), has no target variable or desired output. The goal of unsupervised learning is to infer the natural structure present within a set of data points. Input data is analyzed and grouped together based on the proximity of input values to each other. The groups are then segmented and labeled. Unsupervised learning is useful for exploratory analysis and for dimensionality reduction. Ontanon, Montana, and Gonzalez (2014) describe the process as “learning from observation” (LFO) and explain that the process discovers a “mapping” from the perceived state of the environment and actions.

Machine learning has made continued progress in developing methods that can generalize from data, adapt to changing environments, and improve performance with experience, as well as progress in understanding fundamental underlying issues. By integrating over the distribution of opponent strategies rather than taking a simple empirical average, insights from game theory can be used to derive novel learning algorithms (Blum 2008). Applying machine learning to game theory may shed light on possible opponent strategies by improving a program by playing a game many times against a knowledgeable opponent player. Unsupervised learning enables modeling of the real-world and red forces when sensor data is not matching known (supervised learning) constructs. It can support the classification of sensor observations and predict some inferential knowledge. Doherty et. al. (2016) propose using a multi-step process of

first employing unsupervised learning to explore unlabeled datasets to cluster and classify information in order to construct a supervised classification model. Therefore, the AI system is, in a sense, training itself in an automated fashion.

Most game-learning algorithms are designed to improve a program based on watching or playing against knowledgeable opponent players. Although it is certainly important to understand how a program (or player) could learn from good players, it is equally important to know how those good players became good in the first place. Kim (2018) explains that some learning work has considered how programs might become strong players while relying neither on active analysis nor on experience with experts. Most of these approaches can be considered as self-play, in which either a single player or a population of players evolves during competition on large numbers of contests.

F. Game Theoretic Methods

Game theory methods encompass a wide range of behavioral relations among players and is an umbrella term for the science of logical decision-making in humans and computers. Several game theoretic methods can support the naval tactical predictive analytics application. These include descriptive interpretation, normative (or prescriptive) interpretation, counterfactuals, and regret minimization.

Descriptive interpretation is a way of viewing game theory that attempts to predict how an adversary will act and respond in different strategic settings. This ability was included as part of the conceptual PA capability. Descriptive interpretation suggests that game theory can successfully predict how an adversary will make decisions given a set of circumstances. This method assumes that the game players are rational and will act to maximize their payoffs. While this method provides insights, it will be limited by the imperfect knowledge held by both the blue force and the red force.

Normative or prescriptive interpretation is a game theory method of selecting the best COAs for players. It is prescriptive in that this method determines what the player “should do,” rather than actually predicting what a player might do. Normative (prescriptive) interpretation is a fundamental approach to the conceptual PA capability proposed in this paper. This method attempts to determine the best blue force COA based on what is established as the best outcome, or the most desired 3rd order effect. A Nash equilibrium (an important game theory concept) constitutes a player’s best response to the actions of other players. Thus, the conceptual PA capability could provide an analytical way to determine when the blue force’s COA constitutes a Nash equilibrium. However, it is important to note that there are situations in which it is best to play a non-equilibrium strategy if one expects the red force to do so, or if blue force assets need to be conserved for a longer-term mission.

Counterfactuals (another game theory concept) are claims or hypotheses that are contrary to the facts. A counterfactual can be thought of as a hypothetical state of the world used to assess the impact of action. Counterfactuals are often written as conditional statements in which the conditional clause is false—imagining hypothetically what could have happened. Counterfactual thinking can support the conceptual PA capability by considering as many possible hypothetical future states as possible and analyzing them to eliminate undesired COAs.

A related game theory method is regret minimization. This is a method of running many possible counterfactual hypotheses and carefully altering different COA decisions in each run (or game) to see if this has a positive or negative effect on the outcome. Regret refers to how much better a player would have done if they had made one decision over another at a specific decision point in the game.

Zinkevich et. al. (2007) describe counterfactual regret minimization (CFR) as a self-play algorithm that learns to play a game by repeatedly playing against itself. It starts with a strategy that is uniformly random, where it will play every action at every decision point with equal probability. It simulates playing games against itself and after every game, it revisits decisions and finds ways to improve its strategy. It repeats this process for all combinations of games (which can amount to millions or billions of runs), improving its strategy each time. As it plays, it gets closer and closer towards an optimal strategy for the game: a strategy that can do no worse than tie against any opponent. The way it improves over time is by summing the total amount of regret it has for each action at each decision point and selecting the combination with the least amount of regret. Positive regret for a particular COA means that the blue force would have done better if they had taken that action more often. Negative regret means that the blue force would have done better by not taking that action at all.

After each game in CFR with the program playing against itself, it computes and adds in the new regret values for all of the decisions it just made. It then recomputes its strategy so that it takes actions with probabilities proportional to their positive regret. If an action would have been good in the past, then it will choose it more often in the future. It repeats this process for billions of games. Therefore, CFR produces a long sequence of strategies that it was using on each game. Counter-intuitively, that sequence of strategies does not necessarily converge to anything useful. However, if you compute the average strategy over those billions of strategies in the sequence, then that average strategy will converge towards a Nash equilibrium for the game. In a chess-like game, this average strategy can then be used against any opponent.

However, naval tactical situations are vastly more complex than chess-like games. The numbers and types of blue force and red force warfare resources, tactics, and COAs are

not fixed and sequential as they are in a chess game. In naval operations, the resources, tactics, and COAs can be vast, dynamic, changing over time, can occur at any time, and are largely unknown to the opponent. Therefore, the game theory approach must function with incomplete information and large numbers of instances. Zinkevich et. al. (2007) describe a process of implementing regret minimization in games with incomplete information to determine a Nash equilibrium for very large instances to minimize counterfactual regret which minimizes overall regret. In this process, a framework creates an abstraction of a particular decision point to approximate the behavior of the CFR. These approximations are then mapped back into the full game. Brown et. al. (2019) explain that this CFR abstraction method can be manual and domain specific and may miss strategic nuances of the game. They describe the use of Deep CFR which uses deep neural networks instead of the CFR abstraction to approximate the behavior of CFR in the full game. Deep CFR shows promise as a game theoretic method for PA but requires significant computational power.

V. Conclusion

This paper presented a conceptual framework for applying PA to naval tactical decisions as an automated battle management aid. It described the need to develop and maintain knowledge models of the blue force, red force, and operational situation and described how these models are required for a future PA capability. The paper evaluated AI and game theoretic methods, describing how a combination of statistical, graph theory, decision theory, learning-driven, and game theory methods can be applied to enable a future PA capability.

The payoff for implementing a PA capability as part of an automated battle management aid is predicting 1st, 2nd, and 3rd order effects of possible COAs in order to make the most effective tactical decisions. This real-time wargaming capability would enable short-term and long-term objectives to be weighed, contributing to ensuring that preferred outcomes are more likely. A PA capability could bridge the gap between tactical and strategic thinking, emphasizing causal efficacy – or consequences of actions. In the decades ahead, the Navy will need to maintain maritime decision superiority by incorporating strategic thinking into naval tactical decisions – this can be accomplished with predictive analytics.

References

- Abbott, Dean. 2014. *Applied Predictive Analytics: Principles and Techniques for the Professional Data Analyst*. Hoboken, NJ: Wiley & Sons.
- Bilusich, Daniel, Fred Bowden, and Svetoslav Gaidow. 2011. “Applying Influence Diagrams to Support Collective C2 in Multina-

- tional Civil-Military Operations.” Presented at the 16th International Command and Control Research and Technology Symposium.
- Blum, Avrim and Yishay Mansour. 2007. “Learning, Regret Minimization, and Equilibria.” Chapter 4 (pp. 79-102) in book, *Algorithmic Game Theory*, Edited by Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani, Cambridge University Press.
- Brown, Noam, and Tuomas Sandholm. 2018. “Superhuman AI for Head-up No-limit Poker: Libratus Beats Top Professionals.” *Science* Vol 359, pp. 418-424, January 2018.
- Brown, Noam, Adam Lerer, Sam Gross, and Tuomas Sandholm. 2019. “Deep Counterfactual Regret Minimization.” *Proceedings of the 36th International Conference on Machine Learning (ICML)*, Long Beach, CA.
- Campos, Cassio, and Qiang Ji. 2008. “Strategy Selection in Influence Diagrams Using Imprecise Probabilities.” *Proceedings of the 24th Conference on Uncertainty in Artificial Intelligence*. Helsinki, Finland, July 2008.
- Cox, Louis. 2015. *Breakthroughs in Decision Science and Risk Analysis*. Hoboken, New Jersey: Wiley & Sons.
- De Campos, Cassio, and Qiang Ji. 2008. “Strategy Selection in Influence Diagrams using Imprecise Probabilities.” Presented at the 24th Conference on Uncertainty in Artificial Intelligence.
- Doherty, Conor, Steven Camina, Kevin White, and Gary Orenstein. 2016. *The Path to Predictive Analytics and Machine Learning*. Sebastopol, CA: O’Reilly.
- Bowling, Michael, Neil Burch, Michael Johanson, and Oskari Tammelin. 2015. “Heads-Up Limit Hold’em Poker is Solved.” *Science* Vol 347, Issue 6218, pp. 145-149. 09 January 2015.
- Kalaian, Sema, and Rafav Kasim. 2017. *Decision Management: Concepts, Methodologies, Tools, and Applications*. Information Resources Management Association (IRMA).
- Kim, Sungwook. 2014. *Game Theory Applications in Network Design*. Information Science Reference.
- Kim, Sungwook. 2018. “Basic Concepts for Game Theory.” Chapter 10 in the book, *Game Theory: Breakthroughs in Research and Practice*. IGI Global: Hershey, PA.
- Kuhn, Max, and Kjell Johnson. 2016. *Applied Predictive Modeling*. New York, New York: Springer.
- Moravcik, Matej, Martin Schmid, Neil Burch, Viliam Lisy, Dustin Morrill, Nolan Bard, Trevor Davis, Kevin Waugh, Michael Johnson, and Michael Bowling. 2017. “DeepStack: Expert-Level Artificial Intelligence in Headd-Up No-Limit Poker.” *Science*, Vol 356, Issue 6337, pp. 508-513. 01 March 2017.
- Ontanon, Santiago, Jose Montana, and Avelino Gonzalez. 2014. “A Dynamic-Bayesian Network Framework for Modeling and Evaluating Learning from Observation.” *Expert Systems with Applications*, Vol 41, Issue 11, pp. 5212-5226, September 2014.
- Pearl, Judea. 2008. “Causal Inference.” *JMLR Workshop and Conference Proceedings*. 6:39-58.
- Russell, Stuart, and Peter Norvig. 2010. *Artificial Intelligence, A Modern Approach*, 3rd edition. Upper Saddle, NJ: Pearson.
- Talbot, Patrick, and Dennis Ellis. 2015. *Applications of Artificial Intelligence for Decision-Making*. Middletown, DE.
- Van Gosliga, Ing Sicco Pier, and Ir Henk Jansen. 2003. “A Bayesian Network for Combat Identification.” Paper presented at the *RTO IST Symposium on Military Data and Information Fusion* and published in RTO-MP-IST-040, Prague, Czech Republic, 20-22 October 2003.
- Zhao, Ying, Anthony Kendall, Bonnie Young, and W. Baumgartner. 2015. “Big Data Architecture and Analytics for Improving Combat Identification and the Common Tactical Air Picture.” Naval Research Program Research Report, Naval Postgraduate School, Monterey, CA.
- Zinkevich, Martin, Michael Johanson, Michael Bowling, and Carmelo Piccione. 2007. “Regret Minimization in Games with Incomplete Information.” Technical Report for Department of Computing Science, University of Alberta.