SYM-AM-23-102

Excerpt from the Proceedings of the Twentieth Annual Acquisition Research Symposium

Acquisition Research: Creating Synergy for Informed Change

May 10–11, 2023

Published: April 30, 2023

The research presented in this report was supported by the Acquisition Research Program at the Naval Postgraduate School.

Acquisition Research Program
Department of defense management
Naval Postgraduate School

Leverage AI to Learn, Optimize, and Wargame (LAILOW) for Strategic Laydown and Dispersal (SLD) of the Operating Forces of the U.S. Navy

Dr. Ying Zhao—is a Research Professor at the Naval Postgraduate School (NPS). Her research focuses on data sciences, machine learning, and artificial intelligence methods for defense military applications such as semantic and social networks, common tactical air pictures, combat identification, logistics, and mission planning. Since joining NPS, Zhao has been a principal investigator (PI) on many projects awarded for Department of Defense (DoD) research projects. She received her PhD in mathematics from MIT and is the co-founder of Quantum Intelligence, Inc. [yzhao@nps.edu]

Dr. Doug MacKinnon—is a Research Associate Professor at the Naval Postgraduate School (NPS). He teaches graduate Operations Research (OR) courses in probability and statistics, wargaming, and simulation. He is also heavily involved in leading data mining research projects that explore domains ranging from Navy acquisition to recruiting and has performed as principal investigator (PI) on many projects awarded for DoD research and has advised 100+ doctoral and master's research studies. He earned his PhD in civil engineering from Stanford University and his undergraduate degree from the U.S. Naval Academy. [djmackin@nps.edu]

Abstract

The Secretary of the Navy disperses Navy forces in a deliberate manner to support Department of Defense (DoD) guidance, policy, and budget. The current strategic, laydown, and dispersal (SLD) process is labor intensive, time intensive, and less capable of becoming agile for considering competing alternative plans. SLD could benefit from the implementation of artificial intelligence. We introduced a relatively new methodology to address these questions which was recently derived from an earlier Office of Naval Research funded project that combined deep analytics of machine learning, optimization, and wargames. This methodology is entitled LAILOW which encompasses Leverage AI to Learn, Optimize, and Wargame (LAILOW). In this paper, we developed a stand-alone set of pseudo data that mimicked the actual, classified data so that experimental excursions could be performed safely. We show LAILOW produces a score from a wargame-like scenario for every available ship that might be moved. The score for each ship increases as fewer resources (e.g., lower cost) are required to fulfill an SLD plan requirement to move that ship to a new homeport. This produces a mathematical model that enables the immediate comparison between competing or alternate ship movement scenarios that might be chosen instead. We envision a more integrated, coherent, and large-scale deep analytics effort leveraging methods that link to existing real data sources to more easily enable the direct comparisons of potential scenarios of platform movement considered through the SLD process. The resulting product could facilitate decision makers' ability to learn, document, and track the reasons for complex decision making of each SLD process and identify potential improvements and efficiencies for force development and force generation.

Keywords: artificial intelligence, machine learning, optimization, strategic laydown, and dispersal, SLD, data mining

Introduction

The Secretary of the Navy (SECNAV) disperses Navy forces in a deliberate manner to support Department of Defense (DoD) guidance, policy and budget. The current strategic, laydown, and dispersal (SLD) process is labor intensive and may be benefited by digitalization, automation and application of AI.

The laydown and dispersal of U.S. Naval forces requires manual manipulation of data via weekly Working Groups, which is manpower intensive, and only presents one option to the Chief of Naval Operations (CNO) and SECNAV for consideration. The current SLD process takes one full year to develop and is not responsive to changes in the operating environment or strategic guidance. For example, there is no mechanism to leverage existing data resources to monitor plan execution and track progress toward completion. The 10 years of projected force laydown optimization problem can be overwhelming. The SLD plan needs more than just simple process revision—it needs a modernization with a holistic design leveraging digitalization, automation, and application of AI.

The objective is to digitalize the SLD process with more automation using a cloud-based SLD database, deep analytics, ML/AI to aid decision making, and reduce manpower requirements to focus on the strategic basis and integration of the SLD Plan for improved efficiency and better-informed decision making.

Literature Review

More specifically, based on a memo from RDML T. R. Williams, former director for Plans, Policy, and Integration (N5) for the Deputy Chief of Naval Operations for Operations, Plans, and Strategy (N3/N5; Williams, 2021), N52 is teaming with industry and academia to modernize the SLD process; the challenges are described in the following phases.

Descriptive Phase

What decisions were made? This phase is focused on developing a new database utilizing modern data analytics to display information in a shareable website. The current SLD database exists on a standalone computer with a single user's access in the Pentagon requiring manual updates. This phase's end state is a cloud-based SLD database accessible to the SLD working group that offers permission controls and features improved analysis and display functions.

Predictive Phase

How are we making decisions? What happens if I make a different decision? This phase's end state is an Excursion Modeling Tool. The goal is to develop a decision support tool that uses existing authoritative data and models SLD excursions to assist in rapid decision making with increased accuracy.

Prescriptive Phase

Are we making the right decisions? This phase's end state will utilize deep analytics including AI to take the SLD calculations and other inputs to evaluate the SLD plan and create an optimized plan by including global and theater posture and time-phased force and deployment data (TPFDD) into the calculations.

Methods

This paper details the methods related to the research questions and the prescriptive phase. We apply a mathematical model (i.e., Leverage AI to Learn, Optimize, and Wargame [LAILOW] model) to address deep analytic aspects of the research. LAILOW was derived from an ONR-funded project that focuses on deep analytics of machine learning, optimization, and wargame, essentially Leveraging AI, and consists of the following steps:

Learn: D data, data mining, machine learning, and predictive algorithms are used to learn the patterns from historical data on what and how decisions were made. Data derived from competing demands refer to the excursion proposals and requirements from fleet commanders, national leaders, and assessment data done in various function areas in different installation locations. The current manual process focuses on balancing the budget of unit moving costs with the known demands. Moving cost is developed from permanent change of station (PCS) orders of manpower and readiness of infrastructure. The data are in the form of structured databases and unstructured data such as PowerPoint slides and .pdf documents.

Optimize: Patterns from learning are represented as Soar reinforcement learning (Soar-RL) rules or AGI Transformer models used to optimize future SLD plans. An SLD plan includes a complete gain or loss of naval assets at each installation, homeport, home base, hub, and shore posture location (Fd) and staff (Fg). The optimization can be overwhelming considering numerous combinations. LAILOW instead uses integrated Soar-RL and coevolutionary algorithms to map a total SLD plan to individual units mentioned in an excursion proposal, assessment report, and other what-if analyses.

Wargame: There might be no or rare data for new warfighting requirements and capabilities. This motivates wargame simulations. An SLD plan can include state variables or problems (e.g., future global and theater posture, threat characteristics), which is only observed, sensed, and cannot be changed. Control variables are solutions (e.g., an SLD plan). LAILOW sets up a wargame between state and control variables. Problems and solutions coevolve based on evolutionary principles of selection, mutation, and crossover.

A LAILOW framework can be set up as a multi-segment wargame played by a self-player and the opponent, as shown in Figure 1. The self-player or defender is the SLD enterprise. The opponent or attacker is the environment including competing demands. When applying LAILOW, we first divide the processes into state variables and decision variables as follows:

State variables: These variables and data can be sensed, observed, and estimated, however, cannot be decided or changed by the self-player. They are the input variables, or problems that the self-player must consider. They are also called *tests* or *attacks* for the SLD enterprise.

Decision variables: These variables are needed to solve the problem using optimization algorithms. In LAILOW, the optimization of the decision variables is achieved by the integration of Soar-RL and coevolutionary search and optimization algorithms (Back, 1996; O'Reilly et al., 2020).

Both opponent (tests) and self-player (solutions) evolve and compete as in a wargame. LAILOW is like a Monte Carlo simulation but guided by ML/AI learned patterns with optimization algorithms. In the wargame, the opponent generates large-scale what-if tests to challenge the self-player to come up with better solutions, e.g., SLD configurations to answer the questions such as "what happens if I choose a different decision?" in a systematic simulation.

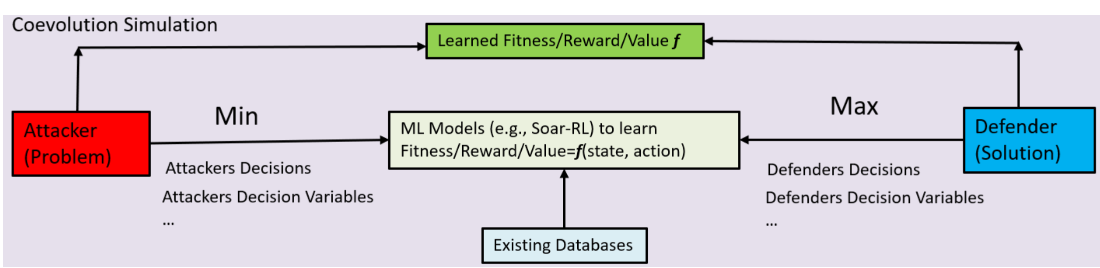

Figure 1. LAILOW viewed in a Coevolutionary wargame simulation; ML algorithms (i.e., Soar-RL) are used to model the fitness or utility functions for both players.

Each "learn, optimize, wargame" cycle dynamically iterates in each stage and across all the value areas with the analytic components and algorithms detailed as follows.

In a LAILOW framework, the “learn” component usually employs supervised ML algorithms such as classification, regression, and predictive algorithms. For example, one can apply a wide range of state-of-the-art supervised ML algorithms from the *scikit-learn* python such as logistic regression, decision trees, *naïve Bayes*, random forest, k-nearest neighbors, and neural networks. Deep learning or AGI Transformers can also be placed in this category where the input data is diversified. An AGI framework typically contains large-scale machine learning models (e.g., billions of parameters in the ChatGPT model; OpenAI, 2023) to learn and recognize patterns from multimodality data.

Supervised ML algorithms can be used to learn the state variables and assessment measures in the function areas for potential SLD and excursion plans such as speed, quality, and fitness of deployment and execution, balance of competing demands and constraints (e.g., avoidance of unacceptable reduction of capability), along with Fd and Fg measures.

In LAILOW, we use Soar-RL to learn two fitness functions separately for the self-player and opponent. In reinforcement learning, an agent takes an action and generates a new state, based on its current state and on the expected value it estimates from its internal model (Sutton & Barto, 2014). It also learns from the reward data from the environment by modifying its internal models. Soar-RL can scalably integrate a rule-based AI system with many other capabilities, including short- and long-term memory (Laird, 2012). Soar-RL carries the following advantages for the military applications, as it

- Can include existing knowledge (e.g., rules of engagements of SLD) and also modify and discover new rules from data
- Learns in an online, real-time, incremental fashion and thus does not require batch processing of (potentially big) data
- Provides the advantage of explainable AI because discovered patterns are represented as rules
- Links to causal learning since it fits the pillars of causal learning (e.g., associations, intervention, and counterfactuals; Pearl & Mackenzie, 2018) by generating the desired effect data using intervention (Wager & Athey, 2018).

The “learn” component can also apply unsupervised learning algorithms. The self-player performs unsupervised machine learning algorithms such as k-means, principle component analysis (PCA), and lexical link analysis (LLA; Zhao & Stevens, 2020; Zhao et al., 2016) for discovering links.

An SLD process needs to perform what-if analysis, as this motivates wargame simulations. An SLD plan can include state variables or problems (e.g., future global and theater posture, threat characteristics, fleet demands to handle these threats), which is only observed, sensed, and cannot be changed. Control variables are solutions (e.g., an SLD plan). LAILOW sets up a wargame between state and control variables. Problems and solutions coevolve based on evolutionary principles of selection, mutation, and crossover.

The number of state and decision variables for an SLD plan and excursion models can be extremely large. Coevolutionary algorithms can simulate dynamic configurations of future warfighting requirements, threats, and global environment and future capabilities, and other competing factors in a wargame simulation. As shown earlier in Figure 1, competitive coevolutionary algorithms are used to solve minmax-problems like those encountered by generative adversarial networks (GANs; Goodfellow et al., 2014; Arora et al., 2017). Adversarial engagements of players can be computationally modeled. Competitive coevolutionary algorithms take a population-based approach to iterate adversarial engagement and can explore a different behavioral space. The use case tests (adversarial attacker population) are actively or passively thwarting the effectiveness of the problem solution (defender). The coevolutionary algorithms are used to identify successful, novel, as well as the most effective means of solutions (defenses) against various tests (attacks). In this competitive game, the test (attacker) and solution (defender) strategies can lead to an arms race between the adversaries, both adapting or evolving while pursuing conflicting objectives.

A basic coevolutionary algorithm evolves two populations with a tournament selection and for variation uses such as crossover and mutation. One population comprises tests (attacks) and the other solutions (defenses). In each generation, engagements are formed by pairing attack and defense. The populations are evolved in alternating steps: First, the test population is selected, varied, updated and evaluated against the solutions, and then the solution's population is selected, varied, updated, and evaluated against the tests. Each test--solution pair is dispatched to the engagement component, and the result is used as a part of the fitness for each of them. Fitness is calculated overall from an adversary's engagements.

Each SLD configuration possesses a fitness value which is related to measures that need to optimize, such as force development (Fd) and force generation (Fg) efficiency. Patterns from "learn" are used to optimize future SLD plans with the measures of the following:

- Cost of an SLD: for a ship to move from one location to another
    - How much does it cost to move personnel: PCS cost per person x # of billets?
    - How much does it cost to prepare requisite infrastructure (matched assessments) to support that ship move?
- Fd/Fg Efficiency: How many excursions or demands are met (matched)?

The optimization can be overwhelming. LAILOW uses integrated Soar-RL and coevolutionary algorithms and simplifies the optimization process.

LAILOW has been used in wargames in DMO and EABO (Zhao, 2021), discover vulnerability and resilience for the logistics operations for Navy ships and Marine's maintenance and supply chain (Zhao & Mata, 2020), and over-the-horizon strike mission planning (Zhao et al., 2020; Zhao & Nagy, 2020).

Use Case

To illustrate the process, we first designed and developed a mock unclassified data set to reflect the SLD process. We began by customizing LAILOW to the SLD process in a high level, as shown in Figure 2. This involved defining self-player variables and opponent variables in the SLD process. Self-player variables are also called defender, control, decision, action, or solution variables. The opponent variables are also called attacker, state, problem, or test variables. Opponent variables include profile variables for a ship such as age, maintenance status, decommission schedule, current installation location, capabilities required by fleet commanders as reflected in the excursion plans, and assessments reflected in a collection of warfighting function areas; these variables are considered pre-determined and known information for a ship and cannot be easily changed for decision makers (defenders) at the time of the SLD process. Attacker variables are the state variables for the defenders to address. Decision variables include move (to what location) or stay, cost, manpower, and maintenance readiness, and are also known as defender variables. Both the defenders and attackers evolve and coevolve, and both are guided by their own fitness functions that reflect the self-player and opponent's competing objectives.

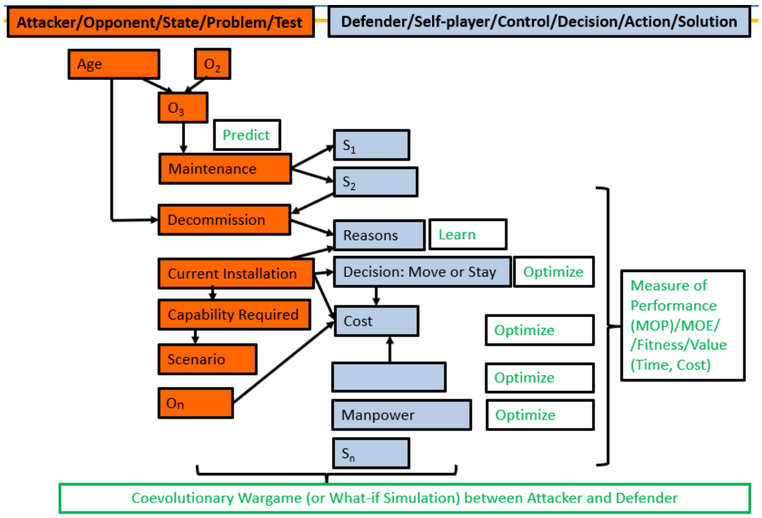

Figure 2. The LAILOW is tailored to the SLD process in a high level to reflect the what-if decision process used by decision makers in the process.

Figure 3 shows the unclassified mock data set to reflect the understanding of the SLD process in Figure 2.

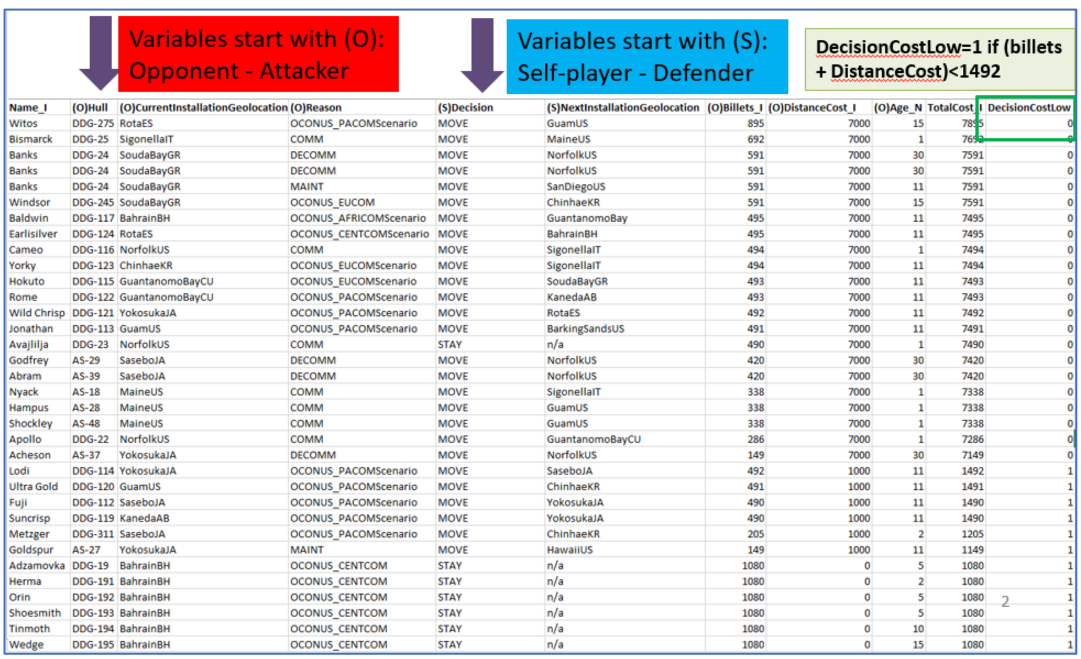

Name_I	(O)Hull	(O)CurrentInstallationGeolocation	(O)Reason	(S)Decision	(S)NextInstallationGeolocation	(O)Billets_I	(O)DistanceCost_I	(O)Age_N	TotalCost_I	DecisionCostLow
Witos	DDG-275	RotaES	OCONUS_PACOMScenario	MOVE	GuamUS	895	7000	15	78[illegible]5	0
Bismarck	DDG-25	SigonellaIT	COMM	MOVE	MaineUS	692	7000	1	769[illegible]	0
Banks	DDG-24	SoudaBayGR	DECOMM	MOVE	NorfolkUS	591	7000	30	7591	0
Banks	DDG-24	SoudaBayGR	DECOMM	MOVE	NorfolkUS	591	7000	30	7591	0
Banks	DDG-24	SoudaBayGR	MAINT	MOVE	SanDiegoUS	591	7000	11	7591	0
Windsor	DDG-245	SoudaBayGR	OCONUS_EUCOM	MOVE	ChinhaeKR	591	7000	15	7591	0
Baldwin	DDG-117	BahrainBH	OCONUS_AFRICOMScenario	MOVE	GuantanomoBay	495	7000	11	7495	0
Earlisilver	DDG-124	RotaES	OCONUS_CENTCOMScenario	MOVE	BahrainBH	495	7000	11	7495	0
Cameo	DDG-116	NorfolkUS	COMM	MOVE	SigonellaIT	494	7000	1	7494	0
Yorky	DDG-123	ChinhaeKR	OCONUS_EUCOMScenario	MOVE	SigonellaIT	494	7000	11	7494	0
Hokuto	DDG-115	GuantanomoBayCU	OCONUS_EUCOMScenario	MOVE	SoudaBayGR	493	7000	11	7493	0
Rome	DDG-122	GuantanomoBayCU	OCONUS_PACOMScenario	MOVE	KanedaAB	493	7000	11	7493	0
Wild Chrisp	DDG-121	YokosukaJA	OCONUS_PACOMScenario	MOVE	RotaES	492	7000	11	7492	0
Jonathan	DDG-113	GuamUS	OCONUS_PACOMScenario	MOVE	BarkingSandsUS	491	7000	11	7491	0
Avajlilja	DDG-23	NorfolkUS	COMM	STAY	n/a	490	7000	1	7490	0
Godfrey	AS-29	SaseboJA	DECOMM	MOVE	NorfolkUS	420	7000	30	7420	0
Abram	AS-39	SaseboJA	DECOMM	MOVE	NorfolkUS	420	7000	30	7420	0
Nyack	AS-18	MaineUS	COMM	MOVE	SigonellaIT	338	7000	1	7338	0
Hampus	AS-28	MaineUS	COMM	MOVE	GuamUS	338	7000	1	7338	0
Shockley	AS-48	MaineUS	COMM	MOVE	GuamUS	338	7000	1	7338	0
Apollo	DDG-22	NorfolkUS	COMM	MOVE	GuantanomoBayCU	286	7000	1	7286	0
Acheson	AS-37	YokosukaJA	DECOMM	MOVE	NorfolkUS	149	7000	30	7149	0
Lodi	DDG-114	YokosukaJA	OCONUS_PACOMScenario	MOVE	SaseboJA	492	1000	11	1492	1
Ultra Gold	DDG-120	GuamUS	OCONUS_PACOMScenario	MOVE	ChinhaeKR	491	1000	11	1491	1
Fuji	DDG-112	SaseboJA	OCONUS_PACOMScenario	MOVE	YokosukaJA	490	1000	11	1490	1
Suncrisp	DDG-119	KanedaAB	OCONUS_PACOMScenario	MOVE	YokosukaJA	490	1000	11	1490	1
Metzger	DDG-311	SaseboJA	OCONUS_PACOMScenario	MOVE	ChinhaeKR	205	1000	2	1205	1
Goldspur	AS-27	YokosukaJA	MAINT	MOVE	HawaiiUS	149	1000	11	1149	1
Adzamovka	DDG-19	BahrainBH	OCONUS_CENTCOM	STAY	n/a	1080	0	5	1080	1
Herma	DDG-191	BahrainBH	OCONUS_CENTCOM	STAY	n/a	1080	0	2	1080	1
Orin	DDG-192	BahrainBH	OCONUS_CENTCOM	STAY	n/a	1080	0	5	1080	1
Shoesmith	DDG-193	BahrainBH	OCONUS_CENTCOM	STAY	n/a	1080	0	5	1080	1
Tinmoth	DDG-194	BahrainBH	OCONUS_CENTCOM	STAY	n/a	1080	0	10	1080	1
Wedge	DDG-195	BahrainBH	OCONUS_CENTCOM	STAY	n/a	1080	0	15	1080	1

Figure 3. An unclassified data set designed and developed to reflect the understanding in Figure 2

Results

We input the mock data into the LAILOW software and simulate a large number of alternative configurations of Navy assets using the mock data set. Figure 4 shows LAILOW solutions as heatmaps (solutions). Each cell in each iteration (i.e., generation in the coevolution algorithm), e.g., circled as 1, 2, and 3, represents a potential SLD plan (Defender) against an environmental test (Attacker), is produced. The heat color shows the fitness for the solution. Clicking on the heatmap cell shows the detail of the corresponding solution configuration.

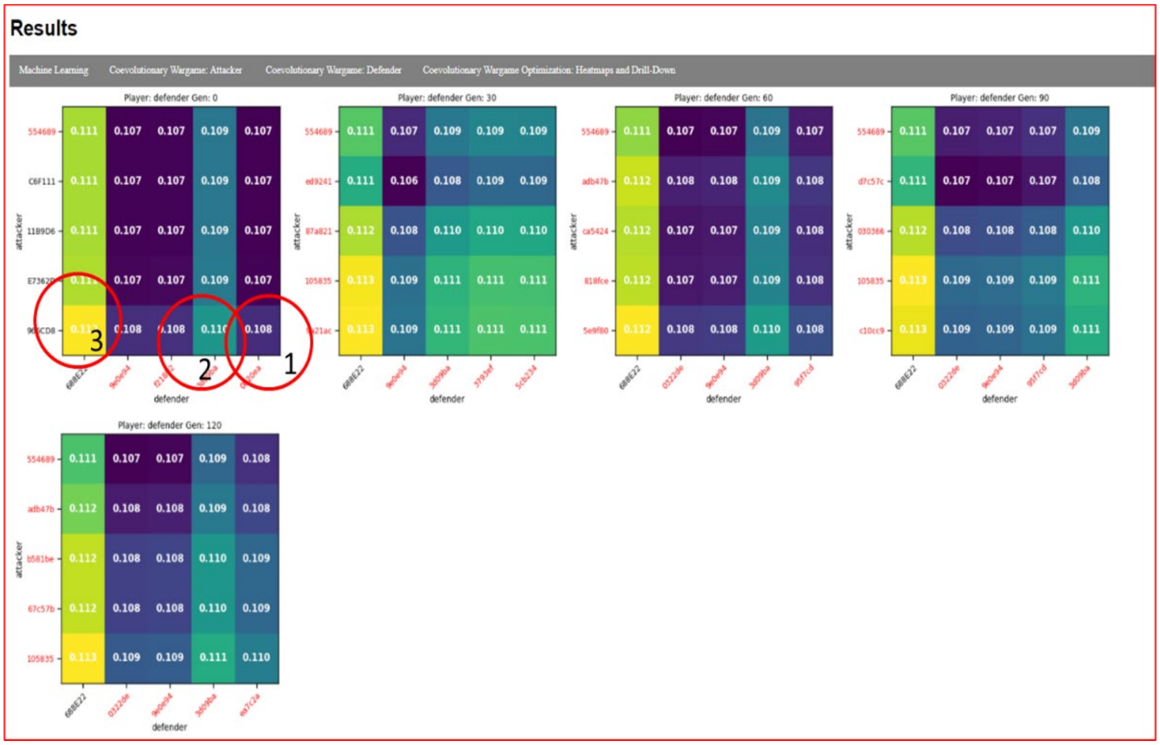

Figure 4. LAILOW solutions as heatmaps (solutions)

We drill down details of the LAILOW simulation in Figure 4. As shown in Figure 5. The LAILOW software illustrates that better decision configurations (6) than ones in the historical databases (4 and 5) can be discovered using the LAILOW software.

Sequence	Variable	Variable Name	Mean	Soar-RL_1_1	Soar-RL_0_1	Soar-RL_1_0	Soar-RL_0_0				Defender's Reward
0	F0	(O)Age_bt_02_08	0.467889908	0.000605917	0.000137289	-1.90E-06	1.27E-06	1			
9	F9	(O)CurrentInstallationGeolocation_KanedaAB	0.009174312	1.20E-05	0.000731227	0	-6.33E-07	1			
41	F41	(O)Hull_DDG-119	0.009174312	1.20E-05	0.000731227	0	-6.33E-07	1		4	
128	F128	(O)Reason_OCONUS_PACOMScenario	0.018348624	-1.36E-06	0.000744571	0	-6.33E-07	1			
138	F138	(S)Decision_MOVE	0.256880734	-0.00019516	0.00093837	1.27E-06	-1.90E-06	1			
155	F155	(S)NextInstallationGeolocation_YokosukaJA	0.018348624	2.40E-05	0.000719252	0	-6.33E-07	1			0.108772543

Sequence	Variable	Variable Name	Mean	Soar-RL_1_1	Soar-RL_0_1	Soar-RL_1_0	Soar-RL_0_0				Defender's Reward
2	F2	(O)Age_lt_02	0.091743119	-5.75E-05	0.000800722	5.93E-11	-6.33E-07	1			
10	F10	(O)CurrentInstallationGeolocation_MaineUS	0.055045872	-4.09E-06	0.000747293	5.93E-11	-6.33E-07	1			
41	F41	(O)Hull_DDG-119	0.009174312	1.20E-05	0.000731227	0	-6.33E-07	1		5	
123	F123	(O)Reason_COMM	0.073394495	-8.14E-05	0.000824637	5.93E-11	-6.33E-07	1			
138	F138	(S)Decision_MOVE	0.256880734	-0.00019516	0.00093837	1.27E-06	-1.90E-06	1			
142	F142	(S)NextInstallationGeolocation_ChinhaeKR	0.027522936	1.05E-05	0.000732682	0	-6.33E-07	1			0.107222755

Sequence	Variable	Variable Name	Mean	Soar-RL_1_1	Soar-RL_0_1	Soar-RL_1_0	Soar-RL_0_0				Defender's Reward
0	F0	(O)Age_bt_02_08	0.467889908	0.000605917	0.000137289	-1.90E-06	1.27E-06	1			
6	F6	(O)CurrentInstallationGeolocation_GuamUS	0.073394495	7.00E-05	0.000673199	0	-6.33E-07	1		6	
41	F41	(O)Hull_DDG-119	0.009174312	1.20E-05	0.000731227	0	-6.33E-07	1			
128	F128	(O)Reason_OCONUS_PACOMScenario	0.018348624	-1.36E-06	0.000744571	0	-6.33E-07	1			
138	F138	(S)Decision_MOVE	0.256880734	-0.00019516	0.00093837	1.27E-06	-1.90E-06	1			
142	F142	(S)NextInstallationGeolocation_ChinhaeKR	0.027522936	1.05E-05	0.000732682	0	-6.33E-07	1			0.108861738

4 is the original in the database, 6 is better than 4 and 5 (in terms of lower cost)

Figure 5. Better decision configurations (6) than ones in the historical databases (4 and 5) can be discovered using the LAILOW software. This shows the potential to discover alternative SLD plans for Naval assets. A defender is an SLD plan for a ship.

Discussions

In reality, the Navy may need to consider many more variables, such as

- Availability of maintenance, pier space, required training schools, etc.
- The policy that requires each ship overseas to return to the continental United States within 10 years
- How each unit fulfills tactical and strategic requirements that must be maintained
- Unseen political pressures that can outweigh numerically based resource requirements

We anticipate our findings to guide the way forward toward further exploration in this area through our suggested methodology. This would likely save time and energy of the decision makers and offer otherwise undiscovered potential alternative solutions toward the development of future SLD plans. In consideration of future efforts, we envision a more integrated, coherent, and large-scale deep analytics effort leveraging methods that link to existing data sources to enable direct comparisons of potential scenarios of platform movement considered through the SLD process. The resulting product could facilitate decision makers' ability to learn, document, and track the reasons for complex decision making of each SLD process and identify potential improvements and efficiencies.

Conclusions

We demonstrate the feasibility of the methodologies of leveraging AI to learn, optimize, and wargame (LAILOW) using mock data. LAILOW produces a score derived from a wargame-like scenario for every available ship that might be moved to a new homeport. The score for each ship increases as fewer resources (i.e., lower cost) are required to fulfill an SLD plan requirement to move that ship to a new homeport. This produces a mathematically optimal response and enables an immediate comparison between competing or alternate ship movement scenarios that might also be chosen, thus improving the automation, consistency, and efficiency of the SLD process.

Acknowledgment

We thank the Naval Postgraduate School (NPS) Naval Research Program (NRP) and the Office of Naval Research (ONR) Naval Enterprise Partnership Teaming with the Universities for National Excellence (NEPTUNE 2.0) program for supporting this research. We would also like to thank the OPNAV N3/N5 team for sponsoring the research topic and providing insightful discussions.

Disclaimer

The views presented are those of the authors and do not necessarily represent the views of the U.S. Government, DoD, or their Components.

References

Arora, et al. (2017). Generalization and equilibrium in generative adversarial nets (GANS). In *International Conference on Machine Learning*, PMLR, 224–232.

Back, T. (1996). *Evolutionary algorithms in theory and practice: Evolution strategies, evolutionary programming, genetic algorithms*. Oxford University Press.

Goodfellow, et al. (2014). *Generative adversarial networks*. arXiv preprint arXiv:1406.2661.

Laird, J. E. (2012). *The Soar cognitive architecture*. MIT Press.

OpenAI. (2023). https://openai.com

O'Reilly, U., et al. (2020). Adversarial genetic programming for cyber security: A rising application domain where GP matters. *Genetic Programming and Evolvable Machines, 21*(1), 219–250.

Pearl, J., & Mackenzie, D. (2018). *The book of why: The new science of cause and effect*.

Sutton, R. S., & Barto, A. G. (2014). *Reinforcement learning: An introduction*. MIT Press.
Wager, S., & Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, *113*(523), 1228–1242.
Williams, T. R. (2021). *Modernizing the strategic lay down and dispersal process* (N3/N5, unclassified). [Email communication with the N3/N5 team on 5/14/2021].
Zhao, Y. (2021). Developing a threat and capability coevolutionary matrix (TCCW) - Application to shaping flexible C2 organizational structure for distributed maritime operations (DMO). In *The 18th Annual Acquisition Research Symposium*.
Zhao, Y., MacKinnon, D. J., Gallup, S. P., & Billingsley, J. L. (2016). Leveraging Lexical Link Analysis (LLA) to discover new knowledge. *Military Cyber Affairs*, *2*(1), 3.
Zhao, Y., & Mata, G. (2020). Leverage artificial intelligence to learn, optimize, and win (LAILOW) for the marine maintenance and supply complex system. In *The 2020 International Symposium on Foundations and Applications of Big Data Analytics* (FAB 2020) in conjunction with the IEEE/ACM ASONAM. https://ieeexplore.ieee.org/document/9381319
Zhao, Y., & Nagy, B. (2020). Modeling a multi-segment war game leveraging machine intelligence with EVE structures. In *Proc. SPIE 11413, Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications II*, 114131V. https://doi.org/10.1117/12.2561855
Zhao, Y., Nagy, B., Kendall, T., & Schwamm, R. (2020). Causal learning in modeling multi-segment war game leveraging machine intelligence with EVE structures. In *Proceedings of AAAI Symposium on the 2nd Workshop on Deep Models and Artificial Intelligence for Defense Applications: Potentials, Theories, Practices, Tools, and Risks*. http://ceur-ws.org/Vol-2819/session3paper3.pdf
Zhao, Y., & Stevens, E. (2020). Using lexical link analysis (LLA) as a tool to analyze a complex system and improve sustainment. In *Unifying themes in complex systems X*. Springer.

Acquisition Research Program
Department of Defense Management
Naval Postgraduate School
555 Dyer Road, Ingersoll Hall
Monterey, CA 93943

www.acquisitionresearch.net