Naval Missile Strike Layout: A Hybrid LLM Agent Architecture for Naval Firepower Planning

1st Yibo Chen
War Research Institute
PLA Academy of Military Science
Beijing, China
1606082029@stu.sqxy.edu.cn

2nd Yang Ping
War Research Institute
PLA Academy of Military Science
Beijing, China
ocean.py@163.com

3rd Shuhang Zhou
College of Mechatronic Engineering
North University of China
Taiyuan, China
372551627@qq.com

***Abstract*—To address the challenge of traditional naval firepower planning methods in balancing scheme rationality, diversity, and robustness under complex battlefield uncertainties, this study proposes a Naval Missile Strike Layout (NMSL) framework, integrating Large Language Models (LLMs) with enhanced conventional planning methods, which achieves the generation, evaluation, and optimal selection of naval firepower planning schemes during pre-combat preparation through multi-LLM collaboration. The architecture consists of four core modules: a situation analysis module, decoupled generation module, route planning module, and scheme evaluation module. Experimental results demonstrate that in two typical naval combat scenarios, NMSL achieves comprehensive performance improvements ranging from 39.1% to 74.7% compared with baseline methods, achieving state-of-the-art performance while validating the necessity and effectiveness of modular design. This study presents the first planning-oriented LLM agent in the domain of naval firepower strike capable of generating a complete strike scheme solely from initial conditions. This research provides a novel framework for firepower planning agents and offers universal architectural design implications for intelligent planning systems with various adversarial styles.**

Keywords—Modern Naval Combat, Firepower Planning, Large Language Model, Intelligent Agent, Decision Support System

I. INTRODUCTION

Modern naval combat has entered the era of "multi-domain joint operations," characterized by highly complex, intensely adversarial, and non-linear battlefield environments [1, 2]. As the core link determining victory or defeat in modern naval warfare, pre-combat firepower planning directly dictates the success or failure of tactical-level decisions, which may entail consequences ranging from squandering precious ammunition and missing critical opportunities to catastrophic outcomes, such as ship destruction, loss of life, and complete mission failure [3]. Modern naval firepower planning demands effective, reliable, and diverse solutions under stringent constraints, including target damage efficacy, resource consumption costs, and battlefield feasibility. Concurrently, the popularization of the relay guidance advanced defense technology and the deployment of anti-ship ballistic missiles (ASBM) [4, 5], further elevates the complexity of firepower planning. Against this backdrop, constructing an intelligent pre-combat firepower planning system with autonomous situational analysis, nested multi-constraint decoupling, and self-assessment capabilities has become a central objective in the global intelligent transformation of naval equipment systems [6].

Traditional pre-combat firepower planning methods based on mathematical theories such as linear programming and game theory exhibit logical transparency and perform well in deterministic scenarios. However, they suffer from three critical limitations. First, traditional methods are confined to scenarios with fixed and precisely located targets; in reality, target positions may change and localization may contain errors, leading to insufficient rationality. Second, traditional methods pursue theoretical optimal solutions and often neglect multi-scheme generation; the lack of diversity makes them vulnerable to targeted counter-strategies in actual combat. Third, traditional methods fail to adequately consider battlefield adversariality and uncertainty, overlooking the impacts of practical factors, such as electronic countermeasures [7], resulting in a lack of robustness. In artificial intelligence, although intelligent algorithms partially address the limitations of traditional methods in non-linear and uncertain scenarios, their black-box nature leads to inadequate interpretability of schemes, failing to meet military decision-making requirements. Notably, the potential of Large Language Models (LLMs) in tactical reasoning has begun to emerge; for example, COA-GPT demonstrates multi-step logical chain reasoning in operational plan generation [8]. However, the direct application of LLMs to firepower planning reveals a weakness: their limited mathematical modeling capabilities.

This study focuses on the rationality, diversity, and robustness of naval firepower planning, with a research scope limited to firepower planning scheme generation and evaluation during pre-combat preparation. The proposed Naval Missile Strike Layout (NMSL) agent framework aims to achieve the following: (a) Establish an intelligent planning mechanism that balances resource optimization and temporal coordination under realistic battlefield constraints; (b) Enhance scheme interpretability and diversity through the combination of semantic understanding and mathematical optimization; (c) Implement a closed-loop evaluation process that supports multi-scheme generation and dynamic selection, thereby improving the scientific rigor and flexibility of command decision-making.

The data supporting the findings of this study are available in the GitHub repository at https://github.com/c12630/NM

II. Design Theory of Naval Firepower Planning Agent

To scientifically construct a LLM-based naval firepower planning agent framework based on theories such as firepower requirement calculation, task decoupling, missile route planning, and LLM agent interpretability, we first introduce the current state of naval firepower planning theory and analyze each theory individually. Then we establish requirements for the agent based on naval firepower planning theories. Finally, we discuss the interpretability required for a LLM-based naval firepower planning agent.

The firepower requirement refers to the amount of firepower required to destroy a target. For independent targets, the firepower requirement typically equals the sum of their anti-missile self-defence capability and the number of missiles needed to achieve complete destruction. For mutually protected targets, the firepower requirement generally equals the combined anti-missile defence capability of multiple ships and the number of missiles required for complete destruction. The significance of firepower requirement calculation in firepower planning lies in establishing a missile allocation strategy for each missile type against each target, thereby providing a reference basis for missile quantity distribution in firepower allocation.

Firepower allocation involves decision-making by friendly combat units regarding strike methods against enemy targets based on total firepower availability, firepower requirements, and target value. Firepower allocation is also constrained by factors such as missile range and no-fly zones. Traditional firepower allocation methods include matrix games, advantage functions, and optimization direction vector methods [9, 10]. However, owing to the diversity, adversariality, and uncertainty of enemy and friendly weapon systems in combat scenarios, traditional firepower allocation methods struggle to achieve rapid and accurate optimal allocations [11].

Missile route planning refers to the design and optimization of flight paths to ensure missile target engagement. Its core lies in comprehensively integrating multiple constraints, including the dynamic distribution of enemy air defence systems, no-fly zones, missile range, turning capability, and multi-missile coordination requirements. Traditional missile route planning methods primarily construct missile paths using algorithms based on geometric principles or optimal control theory [12], thereby achieving satisfactory results under sufficient constraint conditions.

The coupling between multi-missile-type firepower allocation and missile route planning primarily manifests as nested decision-making between the Weapon-Target Assignment (WTA) problem [13] at the firepower allocation layer and the Vehicle Routing Problem with Time Windows (VRPTW) [14] at the route scheduling layer. As both fundamental problems belong to NP-hard problems, their combination leads to a solution space suffering from the curse of dimensionality, which constitutes the essential characteristics and challenges of firepower planning. To address this critical challenge, the methodology of decoupling becomes the key solution approach [15].

In firepower planning theory, common decoupling methods include Task Decomposition, Spatial Decoupling, Temporal Decoupling, Target Decoupling, and Hierarchical Optimization [16-18]. However, these methods typically rely on manual experience and rule-based settings, lack the flexibility to adapt to conditional changes in complex scenarios, and are prone to falling into local optima. Recent advancements in LLMs such as GPT-4o and reasoning LLMs have demonstrated exceptional performance in problem-decoupling [19, 20]. By integrating a Retrieval-Augmented Generation (RAG) knowledge base that provides diverse decoupling methods with modular task binding and reasoning-oriented large models, the framework theoretically achieves human-specified decoupling requirements with flexibility, while retaining scene adaptability and generalization capabilities. In addition, the transparent reasoning process ensures interpretability [21].

For firepower planning schemes, requirements and evaluations can be proposed and measured in three dimensions.

(1)Rationality of the Scheme: In firepower planning, the primary consideration is the rationality of the scheme, that is, its ability to achieve operational objectives to a significant extent. In general, the goal of naval firepower planning is to destroy as many high-value surface targets as possible. To quantify the target value, each target is assigned a specific "reward" score. Statistically, the rationality of the firepower planning scheme was verified through combat simulation to evaluate its destruction score. Mechanistically, rationality depends on the capacity of a scheme to achieve the expected effects.

Traditional complex missile routes must also be streamlined in rational design. Excessive route adjustments caused by target mobility can hinder coordinated saturation strikes, leading to overestimated strike capabilities when multiple routes fail to converge.

(2)Diversity of Schemes: Building upon the premise of scheme rationality, firepower planning can be further subdivided based on distinct tactical objectives, which may include pursuing the highest destruction score, prioritizing the elimination of the most critical targets, or maximizing the number of destroyed targets. Different tactical objectives yield multiple firepower-planning schemes that exhibit rationality. Diverse firepower planning schemes prevent adversaries from adopting targeted strategies, which is important in highly adversarial scenarios.

Designing diverse schemes based on tactical objectives aligns more closely with human decision-making logic than purely algorithmic methods, enhancing interpretability and improving acceptability. From an algorithmic perspective, exhaustive exploration of the entire solution space is impractical because of the combinatorial explosion inherent in firepower allocation [22].

(3)Robustness of the Scheme: Battlefields often exhibit strong uncertainties and adversarial interactions, making robustness a critical consideration. In firepower planning, robustness manifests as an adaptive capability to battlefield uncertainties, ensuring that operational objectives remain achievable despite inevitable uncertainties.

The robustness design primarily operates at two levels: firepower allocation and missile route planning. At the firepower allocation level, ammunition allocation redundancy must be ensured for each critical target during the initial strikes. Reserve ammunition strategies under intense adversarial conditions may fail owing to unit losses and contradict rationality design principles; thus, they are excluded. In missile route planning, considerations must address positional changes due to target evasion and data-link interruptions caused by electronic warfare. Mid-course routes designed to avoid interception require extreme caution. Terminal guidance should rely on autonomous homing capabilities rather than data-link guidance, and time-space-sensitive simultaneous-arrival strategies must be abandoned because of their low fault tolerance.

Finally, in operational planning, LLM-based agents must satisfy interpretability requirements to ensure result acceptability [23]. Interpretability demands that agents fully demonstrate end-to-end reasoning logic from input requirements to operational scheme generation. Interpretability also relies on dual guarantees of theoretical verification and human-machine interaction. On one hand, the "generate-verify-evaluate" process performs self-checks based on operational theories. On the other hand, commanders can directly adjust schemes and constraints via natural language instructions, ensuring human-machine collaboration. Special constraints on manual input should be highlighted for clarity.

III. Construction of the Naval Missile Strike Layout Agent Framework

According to naval firepower planning theory, a complete naval firepower planning process involves the following key steps: (a) calculating the firepower requirements of various munitions needed to strike different targets, (b) determining firepower allocation for each target, (c) planning missile routes and launch times based on firepower allocation, and (d) analysis and evaluation of the scheme. Based on this, the NMSL agent employs multi-LLM collaboration to generate, evaluate, and optimally select naval firepower planning schemes during the pre-combat preparation phase, designing a chain-structured framework with four core modules: 1) Situation analysis module: interprets naval combat situations through DeepSeek R1 based on user natural language input, completing preliminary analysis of target firepower requirements; 2) decoupled generation module: generates three differentiated firepower allocation schemes through DeepSeek R1 reasoning based on different tactical objectives; 3) Route planning module: initializes missile route constraints and invokes algorithms to complete calculations according to firepower allocation schemes and robustness requirements; 4) scheme evaluation module: analyzes and evaluates three complete strike schemes through DeepSeek R1 for decision-making reference.

The NMSL agent framework is illustrated in Fig. 1. Because NMSL is an LLM-based agent, the user's original input can be an unrestricted-format natural language. Effective planning can be achieved as long as the input contains information about the distribution of the enemy and friendly forces, with more detailed inputs yielding better planning outcomes. Upon entering the firepower planning process, it sequentially invokes the situation analysis module, decoupled generation module, route planning module, and scheme evaluation module, ultimately outputting the complete schemes and evaluation results. By utilizing the natural language processing capabilities of LLMs and ensuring seamless data transmission between modules, the system preserves the user's original input across most modules' processing stages. This design allows the agent to directly address specific user constraints and requirements without information loss.

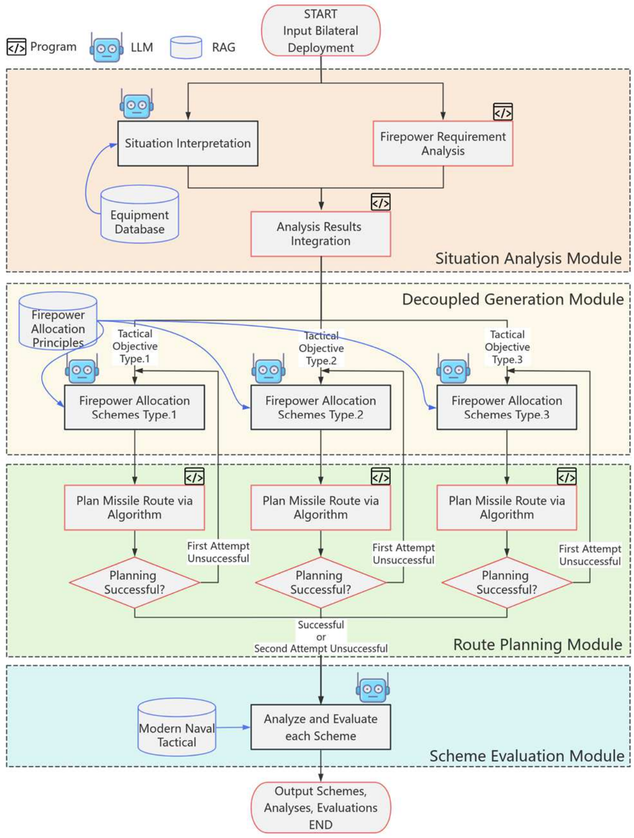

Fig. 1. NMSL Agent Framework

A. *Situation Analysis Module*

The situation analysis module provides a reference basis for formulating firepower allocation schemes consisting of two components: situation interpretation and firepower requirement analysis.

The Situation Interpretation component leverages DeepSeek R1 to analyze the force comparison between opposing parties, providing an intuitive qualitative force assessment. Based on force distribution and anti-ship weapon ranges, it generates a target list within engagement range for each combat unit, establishing a feasible foundation for subsequent firepower allocation. To ensure accurate interpretation by DeepSeek R1, we employ the RAG method to integrate an equipment database. This database automatically supplements detailed performance parameters based on equipment models provided in the user's original input.

The firepower requirement analysis component calculates the munition quantities required for each anti-ship weapon to penetrate target defenses and sink targets according to firepower requirement analysis methods, providing a quantitative basis for ammunition demand in subsequent firepower allocation. In NMSL, firepower requirements are determined by the program following the methodology outlined in Reference [24] by iteratively solving (1) and (2) simultaneously:

$$P = (1\text{-}P_{int})^{M/n} \quad (1)$$

$$n \leq \lg(1\text{-}\sigma) / \lg(1\text{-}P/\omega) \quad (2)$$

Where P is the penetration probability of a single anti-ship missile, P_{int} is the hit probability of the interceptor missile, M is the number of interceptor missiles, n is the number of anti-ship missiles used, σ is the required destruction probability, and ω is the required number of anti-ship missiles that penetrate the defense. The firepower requirement analysis output is presented in tabular form, providing concise munition quantity references for the subsequent firepower allocation.

Finally, the situation analysis module organizes the original user input, outputs from the situation interpretation component, and outputs from the firepower requirement analysis component into complete sentences, which are transmitted to the decoupled generation module.

B. *Decoupled Generation Module*

The decoupled generation module generates target engagement sequences and firepower allocation for the route planning module based on the original user input, situational interpretation, and firepower requirement analysis information.

Regarding the firepower allocation principles integrated into the decoupled generation module, we provide the following knowledge: How to determine "air defense communities" based on each target's interception range and relay guidance capability, then calculate the ammunition consumption required to penetrate these "air defense communities" based on firepower requirements. Building upon this framework, users can autonomously add restrictive rules, combat doctrines, and weapon employment doctrines through original input to ensure human-computer interaction.

Following firepower allocation principles, the module takes the original user input and situational analysis results as inputs using three distinct tactical purposes as decoupling points. Three parallel DeepSeek R1 instances were each bound to a tactical purpose for reasoning analysis, producing firepower allocation schemes based on different tactical objectives. The three tactical purpose prompts are: (a) prioritizing the elimination of high-value targets by concentrating attacks from a specific direction to breach enemy ship defenses, create an access channel, and subsequently destroy high-value targets; (b) maximizing target elimination quantity by prioritizing vulnerable enemy ships; and (c) purchasing maximum score through mathematical expectation calculations combining destruction probability and target value, selecting schemes with higher expected scores. We require the decoupled generation module to consider preliminary constraints based on the maximum range of missiles when generating strike allocations [25], which preserves the planning space to the greatest extent while effectively preventing subsequent route planning failures.

Finally, the module outputs three tactical-purpose firepower allocation schemes to the route planning module in a standardized format.

C. *Route Planning Module*

The route planning module contains three parallel route planners corresponding to the three outputs from the decoupled generation module. Each planner employs heuristic algorithms that take formatted firepower allocation schemes as input, using missile speed, combat unit-target assignments from firepower allocation schemes, target engagement sequences, and robustness requirements including "latest launch time" and "target ambiguous position range" as constraints to determine specific weapon routes and launch timings. To adapt to the input of the decoupled generation module and meet the rationality and robustness requirements of NMSL, we specifically designed a genetic algorithm named multi-target missile routing genetic algorithm (MTMRGA) to find a feasible solution. Compared to conventional missile route planning algorithms, MTMRGA introduces three new constraints: target engagement sequence, target ambiguous position range, and latest weapon launch time, while removing the solving process for firepower allocation, as firepower allocation schemes are directly provided by the decoupled generation module. In the MTMRGA, each individual in the population represents a set of possible route plans. The algorithm iteratively evolves through selection, crossover, and mutation operations, evaluated by a comprehensive fitness function that includes the number of turns and voyage range. It eventually converges to a route planning solution that satisfies all constraints while achieving relatively optimal performance. MTMRGA uses fixed-length real-valued encoding, a population size of 100, crossover probability of 0.8, mutation probability of 0.2, and terminates when the maximum generation of 200 is reached, stagnation lasts for 50 generations, or diversity falls below 0.01, and handles various constraints via a progressive penalty function.

Finally, if the planning succeeds under constraints, the route planning module formats and outputs the results. If planning fails under constraints, it returns the failure result to the decoupled generation module to regenerate firepower allocation schemes. To prevent infinite loops, after two consecutive planning failures, the system concludes the tactical objective infeasibility and outputs this conclusion to the scheme evaluation module.

D. *Scheme Evaluation Module*

The scheme evaluation module primarily enhances the agent's interpretability.

The module first comprehensively presents the complete reasoning process from the original user input to generated firepower strike schemes, explaining special constraint requirements and potential bias sources, verifying constraint violations, and the existence of hallucinations. Subsequently, based on the original user input and modern naval tactical knowledge, the scheme evaluation module individually evaluates the three schemes provided by the route planning module, assigning a 10-point scale score to each scheme and ranking them from best to worst.

IV. COMBAT SIMULATION VERIFICATION

A. *Experimental Setup*

1) Simulation Platform and Scenario Configuration: Experiments were conducted using the commercial simulation software *Mozi Joint Operations Wargaming System* and *Mozi AI Platform* [26, 27]. This system focuses on modern naval and aerial combat simulations, which is a commonly used combat simulation agent experiment platform. To facilitate experimental verification, the knowledge bases used for NMSL in this study were derived from the *Mozi Database* and the *Mozi User Manual*. Two high-quality public scenarios containing typical naval combat content from the system were selected: the "Butter Dispute" and "Thousand Birds Conflict" [28, 29]. The key scenario information is as follows:

(a) Butter Dispute: A typical counter-centralized battle group scenario set near a gulf. The anti-ship force (red) possesses dispersed defensive installations to counter the offensive of the fleet force's (blue) centralized battle group. Anti-ship forces include three ASBM launch units, four bombers, and one cruiser. The fleet forces consisted of one flagship (650 points), two cruisers (100 points each), and one destroyer (110 points). See Fig. 2 (above).

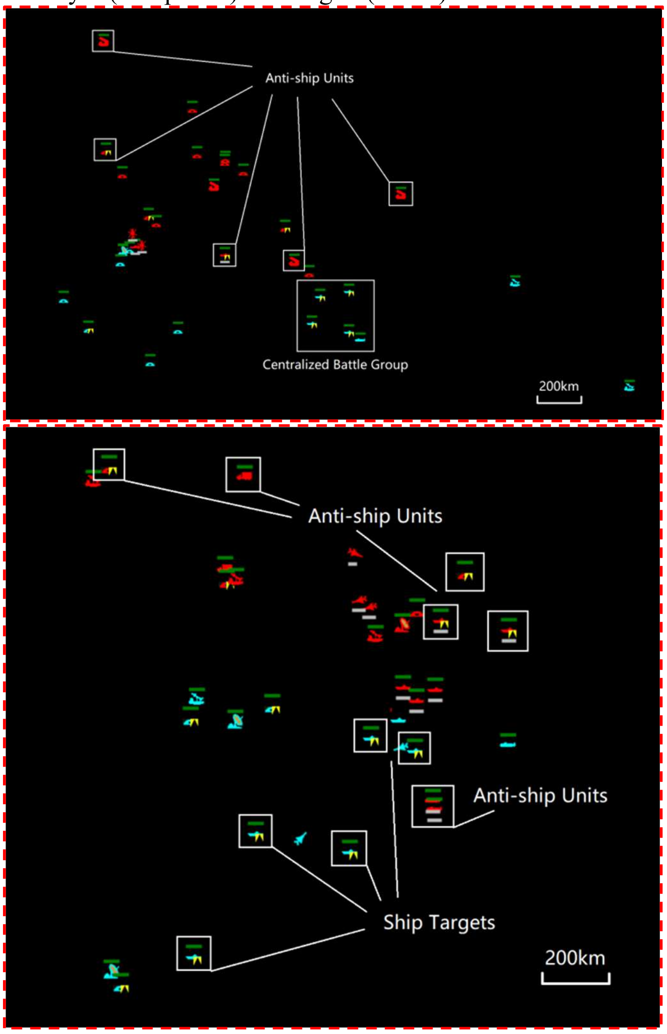

Fig. 2. Scenario Butter Dispute (above) and Thousand Birds Conflict (below) Overview

(b) Thousand bird conflict: A typical distributed naval combat scenario set near a group of islands. The fleet force (blue) possesses dispersed combat units to counter the anti-ship force's (red) maritime offensive. Anti-ship forces include two land-based anti-ship missile launch units, two bombers, one cruiser, one frigate, and one submarine. The fleet forces comprised one flagship (450 points), one cruiser (280 points), and three destroyers (270, 230, and 210 points, respectively). See Fig. 2 (below).

All scenarios included standard pre-deployment mission plans on both sides. To test the robustness, we created rigorous perturbed versions simulating compound disturbances under real combat conditions:

(1) The initial positions of all enemy targets were randomly generated within a 100 km×100 km rectangular area centered on the original coordinates, simulating the opponent's maritime formation changes and electronic warfare-induced positional ambiguity.

(2) None of the missiles can alter their routes after launch, simulating communication jamming effects.

(3) All Red naval vessels are eliminated 20 min after the start of the scenario, simulating combat-induced attrition of friendly units.

2) Evaluation Metric System: The evaluation metrics must measure the rationality and robustness of plans in specific scenarios. The most direct approach for rationality assessment utilizes the target score (Destruction Score S) provided by the scenario, which exclusively considers points from the destroyed surface targets. Note that specific scores vary across scenarios and score comparisons are meaningful only within the same scenario. For intuitive interpretation, the number of destroyed surface targets (Destruction Count N) is also listed. For the robustness evaluation, the ratio between the perturbed version destruction score (Perturbed Versions Score S_R, which could be considered as the comprehensive score of this experiment) and the original destruction score (Robustness Coefficient or Score Retention Rate R) was adopted.

To verify the alignment between the scheme evaluation module's criteria and combat simulation results, the agent's self-assessment score (S_S, an integer from 0 to 10, comparable within the same scenario) was additionally tracked during diversity testing.

B. Baselines

The experiment selected the standard pre-deployed mission plans of the scenarios as a baseline. These plans (generated by *Mozi Joint Operations Wargaming System*) can be regarded as the simplest firepower-planning scheme for each scenario.

We adopted the state-of-the-art (SOTA) method for traditional firepower planning, the genetic algorithm TCGA, as one of the baselines, serving as a representative approach to traditional firepower planning [30].

Because no off-the-shelf LLM-based naval firepower planning agents exist, we endeavored to select representative LLMs. As LLMs themselves serve as highly generalized general-purpose QA tools capable of independent input-output operations, we included LLMs as another experimental baseline. We reinterpreted them by designing input-output interfaces to enable operation within the *Mozi AI Platform* and compatibility with our experiments. We selected DeepSeek R1, a single deep-thinking model, to perform planning through reasoning without modular partitioning and the powerful GPT-4o for direct planning. The temperature values of these LLMs were set to 0.1 and used the same RAG as NMSL.

During testing, the agent inputs included adversarial force distribution information covering the combat unit types, weaponry, positions, target types, weaponry, and possible locations. When multiple scheme outputs were generated, the agent-selected optimal scheme was deployed in the *Mozi Joint Operations Wargaming System*.

Note that we did not directly compare our method with prior methods that take combat simulation platform parameters as inputs and output low-level control commands in real time [31, 32], as such comparisons would be unfair. Our work focuses on exploring how LLMs can generate practical naval firepower planning schemes with diversity and robustness to assist human commanders' decision -making, rather than solving intelligent adversarial gaming problems. Mathematical optimization-based LLMs were excluded from the baselines because preliminary tests revealed their difficulty in understanding constraint mechanisms within naval scenarios and translating practical problems into mathematical formulations. Additionally, NMSL operates independently and can be integrated with methods for intelligent optimization of given schemes [33] and certain dynamic control agents.

C. Test Results

We tested each method in each scenario and its perturbed versions through 100 Monte Carlo experiments and obtained results based on the evaluation metric system, as shown in Table 1.

TABLE I. BASELINE TEST RESULTS

Scenario	Method	S↑	N↑	R↑	S_R↑
Butter Dispute	Pre-deployed	42.2	0.40*	59.2%	25.0
	TCGA*	100.0*	1.00	75%	75.0*
	GPT-4o	0	0	-	0
	DeepSeek R1	208.0	0.32	69.5%*	144.5
	NMSL (Ours)	**231.0**	**2.21**	**87.0%**	**201**
Thou. Conflict	Pre-deployed	239.9	1.13	80.5%	193.1
	TCGA*	480.0*	2.00*	96.5%	463.2*
	GPT-4o	270.0	1.00	**100.0%**	270.0
	DeepSeek R1	691.6	2.92	76.9%	532.1
	NMSL (Ours)	**1015.4**	**3.62**	91.5%*	**929.5**

Bold indicates highest scores, underline denotes second place, * marks third place, and other entries are unmarked.

The experimental results showed that the traditional firepower planning method TCGA achieved significantly better comprehensive scores (perturbed version scores) than the pre-deployed mission and GPT-4o. TCGA does not violate the range constraints; however, the firepower in its generated schemes is overly concentrated, causing excessive damage to a small number of targets. This is because the traditional firepower planning method struggles to "discover" more effective strike strategies.

We observed that GPT-4o's planning often exceeds the range limitations, rendering large portions of its schemes unimplementable, likely because of insufficient coordinate distance calculations. Additionally, GPT-4o's generated schemes exhibit excessive firepower dispersion, leading to inferior performance, sometimes even worse than preset missions (in the Butter Dispute scenario, it failed to destroy any targets). DeepSeek R1 dedicated significant computational effort to distance verification, ensuring the feasibility of the scheme. Its concentrated firepower allocation enables significantly better performance than preset missions and GPT-4o, demonstrating that reasoning LLMs generally outperform conventional LLMs in firepower planning.

Compared with DeepSeek R1, NMSL not only avoids constraint violations but also achieves more rational firepower allocation and strategic planning. Consequently, it demonstrates superior destruction scores and robustness, with destruction score improvements of 11.1% and 46.8% in the two scenarios (relatively), robustness improvements of 17.5% and 14.6% (absolute), and comprehensive scores (perturbed version scores) improved by 39.1% and 74.7% (relative), respectively.

Regarding strike strategies, the critical capability to preemptively eliminate heavily armed anti-air targets (e.g., using cruise missiles against targets carrying theater interceptors or ballistic missiles against targets with self-defense/over-the-horizon interceptors) based on missile characteristics and target-air defense profiles is challenging for traditional optimization algorithms. Remarkably, we found that LLMs can independently consider this aspect, whereas NMSL perfectly implements this strategy under guidance from its RAG knowledge base.

D. Ablation Study

In the ablation study, we tested the contributions of each module to the scheme generation, excluding the scheme evaluation module. We removed three design components from the NMSL (situation analysis, decoupled generation, and route planning modules) and investigated their impacts on the final schemes. Each ablated configuration was tested in 100 Monte Carlo experiments per scenario and its perturbed versions; the results are shown in Table 2.

Key findings from the experimental results are as follows:

The situation analysis module handles substantial mathematical computations, freeing up the cognitive capacity of the subsequent LLM modules to address deeper strategic issues. Without this module, the decoupled generation module would likely spend most of the time calculating firepower requirements or violating range constraints because of inadequate distance calculations, resulting in ineffective strike strategies and poor scheme performance (even inferior to DeepSeek R1). This confirms that the planning resources generated by the situation analysis module are essential for NMSL's proper operation of the NMSL.

TABLE II. ABLATION TEST RESULTS

Scenario	Ablation Config	S↑	N↑	R↑	S_R↑
Butter Dispute	w/o Situation*	100.0*	1.00*	74.0%*	74.0*
	w/o Decouple	0	0	-	0
	w/o Route	210.0	2.00	84.3%	177.0
	NMSL (Ours)	**231.0**	**2.21**	**87.0%**	**201.0**
Thou. Conflict	w/o Situation*	210.0*	1.00*	**100.0%**	210.0*
	w/o Decouple	0	0	-	0
	w/o Route	687.0	2.90	90.2%*	619.9
	NMSL (Ours)	**1015.4**	**3.62**	91.5%	**929.5**

Bold indicates highest scores, underline denotes second place, * marks third place, and other entries are unmarked.

The decoupled generation module serves as NMSL's core component for resolving firepower planning coupling issues, and acts as an indispensable bridge between the situation analysis and route planning modules. Without this, the agent cannot generate any firepower allocation schemes, and the route planning module that relies on it cannot receive effective input, making it impossible for the agent to produce any viable schemes. Compared with the baselines, configurations retaining this module demonstrate improved robustness, suggesting that it is the primary contributor to scheme robustness.

The temporal sequencing of missile launches, as determined by the decoupled generation module, is critical for achieving tactical objectives and determining whether anti-ship missiles can systematically breach target defenses as planned. In contrast, owing to uncertainties in the target positions and the possibility of communication jamming, bypass attack strategies and simultaneous impact tactics, which struggle to ensure robustness, are often discarded. This indicates that detailed missile path design requirements are actually not as significant as traditional research claims. The results show that the route planning module primarily refines the schemes to further enhance the destruction scores.

E. Diversity Testing

In diversity testing, we examined the characteristics of the schemes generated based on different tactical objectives. To verify the alignment between the criteria of the scheme evaluation module and combat simulation results, we recorded the NMSL's self-assessment scores. For each tactical objective in each scenario and its perturbed versions, we conducted 100 Monte Carlo simulations. Table 3 presents the results.

TABLE III. DIVERSITY TEST RESULTS

Scen.	Tactical Obj.	S↑	N↑	R↑	S_R↑	S_S↑
Butter Dispute	High-value Targets	210.0	2.00	85.2%	179.0	7
	Elimination Quantity	**231.0**	**2.21**	**86.6%**	**200.0**	**9**
	Maximum Score*	210.0	2.00	84.3%*	177.0*	6*
Thou. Conflict	**High-value Targets**	**1015.4**	**3.62**	91.5%	**929.5**	**9**
	Elimination Quantity	988.6	3.48	91.3%*	902.2	8
	Maximum Score*	689.3*	2.91*	**95.7%**	659.4*	7*

Bold indicates highest scores, underline denotes second place, and * marks third place.

NMSL's self-assessment score rankings generally align with the experimental results, indicating that its scheme evaluation module aligns with combat simulation outcomes and possesses reliable self-assessment capabilities.

From the results, we observe that directly instructing NMSL to "Maximum Score" paradoxically yields lower scores, whereas specifying explicit tactical objectives produces higher scores. This suggests that the ambiguous "Maximum Score" objective may confuse NMSL.

The results imply that in the Butter Dispute, a typical counter-centralized battle group scenario, prioritizing the quantity of destroyed targets represents better tactical decision-making, whereas in the Thousand Birds Conflict—a distributed naval combat scenario—focusing on high-value targets proves more effective. This distinction arises because centralized group targets are densely clustered with concentrated air defenses. This makes it extremely difficult for anti-ship missiles to penetrate their defense and even more difficult to destroy the flagship. Given limited anti-ship resources, focusing on destroying peripheral escort vessels has become effective. Conversely, distributed naval combat features sparse targets with lower individual defense capabilities, making the prioritization of high-value targets more advantageous under limited anti-ship firepower.

V. DISCUSSION

This study focuses on exploring how to enable LLM-based agents to generate practical naval firepower planning schemes, characterized by both diversity and robustness, during the pre-combat preparation phase to assist human commanders' decision-making, rather than replacing human command responsibilities. To avoid the risk of misuse of autonomous weapon systems, this study was strictly limited to a closed, non-real combat simulation environment for training, analysis, and tactical deduction.. Currently, the complete runtime of NMSL is approximately 1100 s, consuming approximately 38k tokens, which meets the requirements for generating plans during the pre-combat preparation phase. In practical use, to meet the rapid decision-making needs of encounter battles, various naval encounter scenarios can be designed in peacetime, and NMSL can automatically compute decisions for various scenarios; after manual inspection and supplementation, these can be stored as contingency plans. Through small-scale pilot applications, we have found that the contingency plans production pipeline adopting NMSL can reduce total planning time by approximately 90%, greatly saving manpower, but it also places higher hardware demands on the computing server (requiring the capability for full-precision concurrent inference of DeepSeek).

This study has currently set three different types of tactical objectives for NMSL. As research further deepens and naval warfare theory further develops, more tactical objectives can be set in the future, and the NMSL framework can be expanded upon, adding more parallel computing processes in the decoupled generation module and route planning module to further enhance diversity (it is also possible to replace different tactical objectives based on the original framework to improve applicability). In practical use, NMSL can also use its diverse tactical objectives to inversely verify the rationality of fleet deployment and fleet formations, providing professional "red team testing."

VI. CONCLUSION

This study proposes and validates NMSL, a hybrid intelligent agent framework integrating LLM semantic reasoning with mathematical programming for modern naval firepower planning. Through a modular architecture, we, for the first time in the domain of naval firepower strike, realize an end-to-end decision-support pipeline in the form of an LLM agent, enabling a seamless flow from situation understanding and task decoupling to route planning and scheme evaluation.

The core innovation lies in its hybrid design: LLMs handle high-level semantic comprehension, tactical decoupling, and scheme generation, while mathematical programming ensures precise constraint satisfaction and optimization. Experimental results demonstrate that NMSL significantly outperforms traditional methods and single-model baselines across typical naval combat scenarios, achieving high score retention rates (87.0% – 91.5%) under perturbed conditions and improving comprehensive scores by 39.1% – 74.7% over single-LLM architectures. Ablation studies further validate the critical contribution of each module.

This study reveals: First, a hybrid architecture combining LLMs and mathematical programming is an effective paradigm for solving planning problems in complex adversarial environments, as it balances semantic flexibility and computational optimality. Second, the closed-loop "generate-verify-evaluate" process enhanced by a RAG knowledge base improves the interpretability and reliability of the agent's output, providing a trustworthy foundation for human-machine collaborative decision-making. Third, the study validates the feasibility of generating schemes decoupled based on tactical objectives, which not only enhances scheme diversity but also aligns decision logic more closely with human commanders.

This study provides an extensible framework and validation paradigm for designing intelligent military decision-support systems. Future research can proceed along three directions: First, extending multimodal input capabilities (e.g., imagery, radar signals) for more comprehensive battlefield situational awareness. Second, expanding planning to encompass emerging combat forces (e.g., drone swarms) and developing corresponding cooperative planning methods. Third, exploring cross-domain collaborative planning (e.g., air force's escort) to support broader joint operational requirements. The hybrid intelligent agent pathway validated in this study also holds significant reference value for other intelligent planning systems involving resource allocation, path planning, and similar adversarial-style problems.

REFERENCES

[1] 1.Yang, J., Xi, Y.: Implementation Plan of U.S. Military Multi-Domain Operations and Helicopter Performance Upgrade Program. Helicopter Technology 2019(4), 60-66 (2019)

[2] 2.Ye, Q., Wang, Q.: U.S. Military Releases Latest 1.5 Version of Multi-Domain Operations Concept. Military Digest 2019(5), 51-54 (2019)

[3] 3.Jin, B., Sun, H.: Review of Firepower Application Methods in Maritime Operations. Command Control & Simulation 45(2), 6 (2023)

[4] 4.Dai, J.: Study on Shifting Accuracy and Success Probability Based on Relay Guidance. 2011 IEEE International Conference on Information Theory and Information Security(ICITIS 2011), Hangzhou, China, 2011

[5] 5.Binnie, J.: Iran successfully tests radar-guided ASBM. Jane's Defence Weekly: IHS Jane's Defence Weekly 2017(11), 20

[6] 6.Liang, R., Wang, C., Ji, X., et al.: Intelligent Planning Process for Joint Fire Strike. Journal of Command and Control 9(4), 478-486 (2023)

[7] 7.Chen, S., Shen, H., Wang, R., et al.: Research on the Relationship Between Predictive Uncertainty and Adversarial Robustness. Journal of Software 2022(2), 524-538 (2022)

[8] 8.Goecks, V.G., Waytowich, N.: COA-GPT: Generative Pre-Trained Transformers for Accelerated Course of Action Development in Military Operations. In: 2024 International Conference on Military Communication and Information Systems (ICMCIS), Koblenz, Germany, 2024 (2024)

[9] 9.Shi, J., Song, Y.: Multi-phase Fire Allocation for Ship Formation Based on Sequential Game Theory. Command Control & Simulation 37(3), 1-4 (2015)

[10] 10.Wang, H., Wang, X.: Fire Allocation Using Second-order Dynamic Programming Algorithm. Ordnance Industry Automation 25(8), 1-2 (2006)

[11] 11.Yao, T., Wang, Y., Dong, Y., et al.: Application of Deep Reinforcement Learning in Operational Mission Planning. Aerodynamic Missile Journal 2020(4), 16-21 (2020)

[12] 12.Liu, Y.: Research on Cooperative Path Planning Methods for Multiple Cruise Missiles. Ph.D. Thesis, Harbin Institute of Technology (2009)

[13] 13.Yang, S., Huang, J., Liu, Y., et al.: Research on Joint Fire WTA Problem Based on Genetic Algorithm. Computer Simulation 29(3), 1-4 (2012)

[14] 14.Barlow, A.W.: Stochastic Duels for Evaluating Ground-Based Anti-Ship Systems. Ph.D. Thesis, Naval Postgraduate School (2022)

[15] 15.Liu, G., An, Z., Lao, S., & Li, W. (2022). Firepower distribution method of anti-ship missile based on coupled path planning. Journal of Systems Engineering and Electronics, 33(4), 1010-1024.

[16] 16.Li, L., Liu, F., Mei, Y.: Terminal Anti-TBM Weapon-Target Matching Optimization and APSO Solution Algorithm. Systems Engineering and Electronics 2013(5) (2013)

[17] 17.Hu, L.: Research on Missile Effectiveness Evaluation and Cooperative Combat Mission Planning Methods. Ph.D. Thesis, Harbin Institute of Technology (2022)

[18] 18.Li, P., Li, C.: Modeling and Algorithm for Weapon-Target Cooperative Fire Allocation. Command Control & Simulation 000(002), 36-40 (2015)

[19] 19.Zhang, Y., Du, L., Cao, D., et al.: An Examination on the Effectiveness of Divide-and-Conquer Prompting in Large Language Models. arXiv (2024)

[20] 20.Wei, J., Wang, X., Schuurmans, D., et al.: Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. arXiv (2023)

[21] 21.Lewis, P., Perez, E., Piktus, A., et al.: Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. arXiv (2021)

[22] 22.Liu, B., Yang, J., Shi, W.: Ballistic Missile Multi-target Fire Allocation Method. Journal of Projectiles, Rockets, Missiles and Guidance 2020(04), 37-40 (2020)

[23] 23.Luo, H., Specia, L.: From Understanding to Utilization: A Survey on Explainability for Large Language Models. arXiv (2024)

[24] 24.Jiang, C., Zeng, Q.: Calculation Method of Anti-ship Missile Launch Quantity in Saturation Attack. Ship Electronic Engineering 2008(03), 54-56 (2008)

[25] 25.Jia, Z., Lu, F., Wang, H.: Multi-platform Cooperative Attack Mission Planning Based on Decoupling Optimization and Circulation APF. Journal of Beijing University of Aeronautics and Astronautics 2020(06), 1142-1150 (2020)

[26] 26.Hua Shu Defense. "Mozi" Joint Operations Wargaming System. http://www.hs-defense.com/nd.jsp?fromColId=105&id=1#_np=105_1597, last accessed 2024/12/01

[27] 27.Hua Shu Defense. Mozi· Joint Operations Agent Development and Training Platform. http://www.hs-defense.com/col.jsp?id=124, last accessed 2024/12/01

[28] 28.CICC. Announcement on the Award Results of the 6th National Wargaming Competition Scenario Collection Activity. http://ciccwargame.com/h-nd-76.html, last accessed 2024/12/01

[29] 29.CICC. Qiandao Bingfeng_2023 National Competition Scenario. http://ciccwargame.com/h-col-101.html, last accessed 2024/12/01

[30] 30.Xiang, Y.: Research on Optimization Model and Algorithm for Target Allocation Problem in Joint Fire Strike. Ph.D. Thesis, Xidian University (2018)

[31] 31.Liu, Z.: Research on Wargaming Model Based on Multi-agent System. Ph.D. Thesis, Huazhong University of Science and Technology (2016)

[32] 32.Chen, L.: Research on Multivariate Time Series Classification Under Strong Adversarial Conditions with Incomplete Information. Ph.D. Thesis, National University of Defense Technology (2022)

[33] Wang, G., Cao, H., Liu, P., et al.: Research on Combat Effectiveness Optimization Method Based on AHP-DQN. Command Control & Simulation 2023(3), 78-86 (2023)