COA-GPT 2.0: An Agentic AI Planning Tool for Accelerating the Military Decision-Making Process

Vinicius G. Goecks*
DEVCOM Army Research Laboratory
Aberdeen Proving Ground, MD
vinicius.goecks@gmail.com

Nicholas Waytowich*
DEVCOM Army Research Laboratory
Aberdeen Proving Ground, MD
nicholas.r.waytowich.civ@army.mil

Ryan Bookman
Stanford University
Stanford, CA

Christian Kenney
Mission Command Battle Lab
Ft. Leavenworth, KS

Abstract

Developing effective Courses of Action (COAs) for large-scale military operations is a resource intensive, time-consuming, and critical aspect of the Military Decision Making Process (MDMP). The current Command and Control (C2) processes are largely manually intensive, slow and face challenges amid accelerating operational tempos, increased complexity, and adversaries capable of quickly exploiting windows of opportunity. To address this, we introduce COA-GPT 2.0, an advanced agentic artificial intelligence (AI) planning assistant that leverages multimodal Large Language Models (LLMs), specialized external tools, and human expertise to rapidly support the C2 planning staff during MDMP, significantly streamlining mission planning. In this report, we outline the current capabilities of COA-GPT 2.0, present military subject matter expert (SME) assessment of its MDMP product outputs, and define the path forward to enable command posts (CPs) with reduced staff that still maintain robust analytical capabilities. Our goal with COA-GPT 2.0 is to accelerate the military planning process while enabling Commanders to retain decision advantage under compressed decision timelines, dynamically adapting operational plans in response to emerging threats and changing battlefield conditions in the complex, multi-domain battlefields of today and the near-future.

1 Introduction

The development of Courses of Action (COA) in military operations represents one of the most critical and complex planning challenges faced by Commanders and their staff at the Battalion, Brigade, Division, Corps, and Theater Army echelons. At these higher command levels, the Military Decision Making Process (MDMP) [11] serves as the doctrinal framework that structures this complex planning effort. The process intentionally incorporates multiple human-centric features like repetition, multi-form communication, and group ideation to ensure Commanders and their staffs develop sufficient situational understanding to maintain unity of effort across subordinate echelons operating in complex, multi-domain battlespaces. The traditional MDMP approach requires extensive staff resources, domain expertise across multiple warfighting functions, and significant time investment to produce viable COA options that balance feasibility, suitability, acceptability, and distinctiveness — all while operating in high-pressure environments where decisions affect thousands of soldiers and can determine mission success or failure across vast operational areas, as well as possibly recommending the employment to critical national defense and strategic deterrent capabilities.

*Equal contribution

Preprint. Under review.

However, emerging battlefield realities are rapidly outpacing the deliberate tempo of traditional MDMP approaches at all echelons. The proliferation of autonomous systems, ubiquitous sensors, artificial intelligence (AI), and advanced data analytics is dramatically accelerating the operational tempo while exponentially increasing the volume and velocity of information. Future multi-domain operations will feature fleeting windows of opportunity that may exist for only minutes rather than hours or days — windows that cannot be exploited through time-intensive planning cycles. Furthermore, sophisticated adversaries employing advanced anti-access/area-denial (A2AD) capabilities will create contested electromagnetic environments where denied, degraded, intermittent and limited (DDIL) communication conditions become the norm, while precision long-range fires increasingly threaten large, static command posts (CPs), driving the need for agile and mobile command nodes.

We address these challenges through COA-GPT 2.0, an agentic AI-enabled planning assistant that accelerates the planning process and COA generation by reasoning over large amounts of data using the latest generative AI capabilities. COA-GPT 1.0 [7], which was our first generation of a generative AI system for rapid COA development using Large Language Models (LLMs), has evolved substantially in its second iteration. Unlike COA-GPT 1.0 which focused on simple COA generation using a single LLM, the current system comprehensively supports all MDMP steps through an integrated multi-agent architecture, generating the full range of doctrinal planning products while maintaining human oversight.

We envision COA-GPT as an intelligent assistant within the Mission Command Information System, operating as an intermediary between data inputs and human staff. By shouldering routine data-processing and document-drafting tasks through multimodal LLMs tuned by military subject matter experts, planners can focus on higher-order analysis and Commander engagement. Inspired by previous approaches [12, 9], we also aim to provide a platform where the science of war is handled by computers, while the art of war remains the domain of human judgment, improving both the speed and quality of decision-making while preserving the essential human elements of military command.

2 Background and Related Work

2.1 Planning Military Operations and The Military Decision-Making Process (MDMP)

The Military Decision-Making Process (MDMP) [11] serves as the doctrinal framework guiding Commanders and staff through complex planning at Battalion, Brigade, Division, Corps, and Theater Army echelons. This deliberate process encompasses seven distinct steps: Receipt of Mission, Mission Analysis, COA Development, COA Analysis, COA Comparison, COA Approval, and Orders Production. Each step involves deliberate collaboration among the Commander and specialized staff sections, demanding meticulous consideration of multiple variables, including terrain, available forces, enemy capabilities, time constraints, and mission objectives to transform strategic guidance into executable operational plans. Modern planning staffs must also reconcile a torrent of higher-headquarters data with evolving operational environment (OE) considerations under increasingly compressed timelines. Simultaneously, the survivability of large, static command posts has diminished in an era of precision long-range fires and advanced ISR capabilities, necessitating more distributed, mobile command nodes with reduced footprint to enhance survivability and reduce signature.

The core challenge for modern Command and Control (C2) is the need to significantly reduce CP footprints while simultaneously increasing the speed of planning and the quality of the resultant plan. Meeting this dual objective strains traditional processes and drives new technological and procedural priorities, including emission reduction, the integration of AI, support for distributed collaboration, and multi-domain visualization. Together, these converging pressures (data overload from HHQ, compressed timelines, CP survivability, and the modern imperative for smaller, faster C2) expose the vulnerabilities of the MDMP itself, a deliberate and often sequential process that originated in a different era of warfare. The strain on this arguably archaic process highlights the urgent need for its modernization to remain effective on the future battlefield.

2.2 AI for Planning Military Operations

Traditional computational approaches initially attempted to address limitations in the MDMP and military planning through knowledge-based and multimodal human-centered systems such as Deep Green [18] and CADET [12, 9]. Despite their innovative use of multimodal interfaces and anticipa-

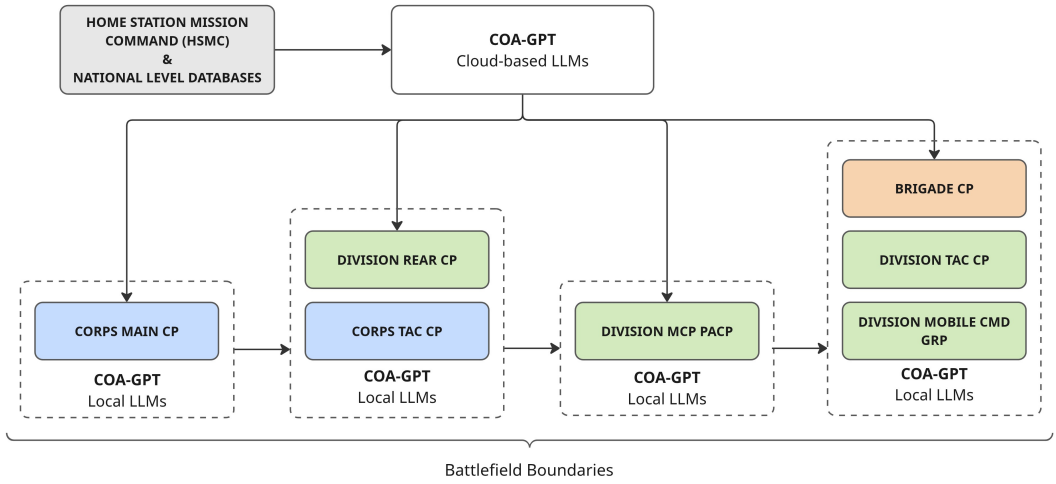

Figure 1: The distributed COA-GPT architecture for military planning, leveraging both local and cloud-based Large Language Models (LLMs) depending on echelon and connectivity available at distinct Command Posts (CPs). The diagram shows integration with Home Station Mission Command (HSMC) and National-Level Databases, as well as different types of command posts: Main Command Post (Main CP), Tactical Command Post (TAC CP), Rear Command Post (Rear CP), Main Command Post / Primary Alternate Command Post (MCP PACP), and Mobile Command Group (CMD GRP).

tory planning techniques, these systems saw limited adoption by military end-users. Reinforcement Learning (RL) also emerged as a promising machine learning-based approach for generating COAs, successfully demonstrating high performance in complex simulations and small-scale tactical scenarios [2, 15, 20, 8]; however, RL-based systems revealed significant drawbacks, including brittle performance in unforeseen situations, computationally intensive training processes, and limited ability to incorporate nuanced human preferences [21]. Other machine learning-based systems focused specifically on optimizing existing COAs through iterative wargaming simulations [17], yet these approaches still relied on human staff or separate computational methods to provide initial COAs.

Recent advances in large language models (LLMs) [1, 3, 5, 14, 19, 10, 13, 4] have markedly improved the integration of human expertise into automated decision-making frameworks, overcoming several limitations associated with previous methods, including the ability to understand more realistic and complex scenarios. In addition, LLMs offer increased ease and efficiency in adapting human domain-specific knowledge, including to military applications [16]. Capitalizing on these advancements, COA-GPT 1.0 [7] demonstrated the practical feasibility of employing LLMs to rapidly generate initial COAs in a small-scale game scenario, effectively integrating multimodal mission data and real-time Commander feedback without retraining the LLM. Building upon the success of COA-GPT 1.0, this report now presents COA-GPT 2.0, which expands beyond merely generating COAs to assisting staff throughout the entire MDMP in a realistic and complete military training scenario. COA-GPT actively supports military planners by automating planning tasks and generating doctrinally-aligned outputs across multiple stages of MDMP.

From here forward, we refer to COA-GPT 2.0 simply as COA-GPT, and we present its capabilities and performance in supporting the military staff throughout the MDMP.

3 COA-GPT

The proposed high-level architecture, shown in Figure 1, is a distributed ecosystem where instances of COA-GPT are deployed across various command echelons. The key architectural difference between these instances lies in the underlying Large Language Model (LLM) technology they employ, which is dictated by their operational environment. At echelons with robust connectivity, COA-GPT can leverage state-of-the-art, cloud-based LLMs. This allows them to interface directly with National Level Databases and perform complex, large-scale data synthesis to support strategic-level planning.

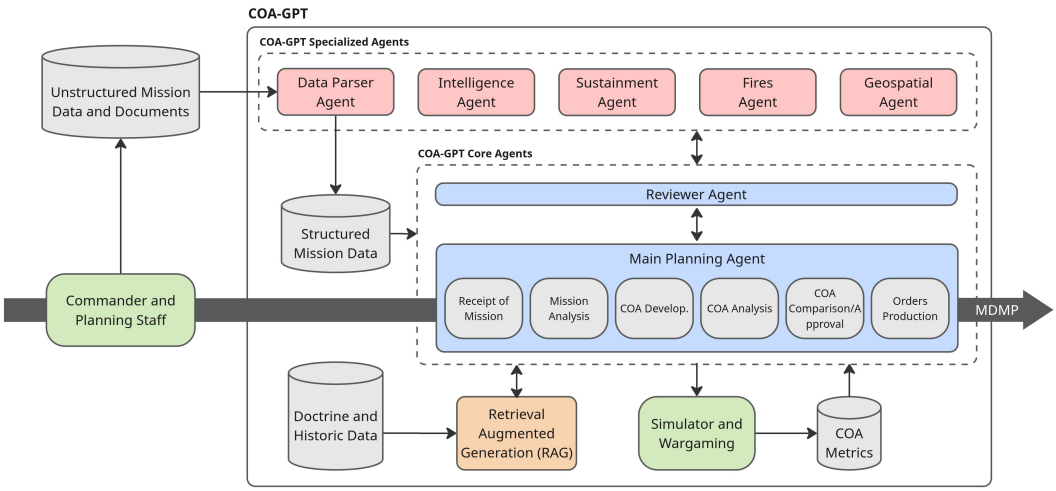

Figure 2: COA-GPT's multi-agent ecosystem for mission planning, showing the interplay between automated planning agents, data analysis, simulation, the planning staff, and the Commander.

Conversely, command echelons deployed within the battlefield operate under the assumption of denied, degraded, intermittent, and limited (DDIL) communications. For these units, which include various command posts (CPs) from the Corps to the Brigade level, connectivity to the cloud cannot be guaranteed, or even desired due to electromagnetic (EM) emission concerns. Therefore, these echelons are equipped with instances of COA-GPT that run on self-contained local LLMs. This critical design choice ensures each command post has resilient, standalone AI assistants that remain fully functional for planning and decision support independent of connectivity.

Figure 2 illustrates each instance of COA-GPT's multi-agent ecosystem, which orchestrates automated planning agents, data analysis, simulation, the planning staff, and the Commander to execute the MDMP. Unlike a single LLM solution, COA-GPT employs multiple specialized agents in a hierarchical, role-specialized organization that mirrors military planning staff functions. The system comprehensively supports all MDMP steps and substeps through this integrated planning architecture, reasoning over large amounts of unstructured data, which is common in military planning, to generate high-fidelity products and recommendations.

The initial phase of mission planning involves ingesting mission-related documentation, which is typically disseminated in various formats such as PDF, Microsoft Office files (e.g., Word, PowerPoint, Excel), GeoJSON, and KML files. The first component of COA-GPT is a specialized Data Parser Agent designed to address this challenge. This agent extracts all relevant information from these distinct sources and transforms the unstructured mission data into a standardized, structured format. This foundational step is essential for ensuring data consistency for subsequent analysis by COA-GPT and both its Core and Specialized agents. During this ingestion process, the system also automatically identifies and flags potential discrepancies within the supplied documents, alerting the planning staff to any inconsistencies.

COA-GPT's knowledge base is also augmented by doctrinal and historical data. To provide doctrinal context, we have developed a military Retrieval-Augmented Generation (RAG) database containing an extensive corpus of publications from the Army Publishing Directorate (APD) and other reputable sources. For historical context, all user feedback and interactions with COA-GPT are stored and leveraged, where applicable, in the planning of future missions. Throughout the planning process, the C2 planning staff or the Commander can also upload additional mission- and scenario-related documents through a graphical user interface (GUI).

The system produces comprehensive planning products throughout the MDMP, including enemy and friendly courses of action, synchronization matrices, warning and operations orders, mission briefs, and other doctrinal outputs. A specialized Geospatial Agent performs terrain analysis from map imagery, while the system maintains running estimates and generates final fighting products for mission execution. The system also performs audio analysis and serves as a comprehensive data management system that maintains a Common Operational Picture (COP) and generates final fighting

products for mission execution. This integrated approach ensures all planning elements remain synchronized while shouldering routine data-processing and document-drafting tasks, allowing human planners to focus on higher-order analysis and commander engagement.

Under this workflow, the planning staff uses COA-GPT to ingest and analyze mission documents and then generate initial drafts for the doctrinal products required at each planning step during MDMP. COA-GPT uses its Main Planning Agent along with RAG on our military dataset to generate initial MDMP outputs and products. The RAG system ensures that these initial products are informed by military doctrine and available mission data. These outputs are then sent to a separate Reviewer Agent whose role is to critique and provide feedback to the main planning agent according to specific criteria informed by doctrine and subject matter experts. The main planning agent uses the feedback provided by the reviewer to generate revised outputs, which in our experience, are significantly improved compared to the original outputs. This iterative workflow can be repeated to improve refinement and produce higher-quality results, though at the cost of additional time and compute resources. During the planning process, the main planning agent periodically tasks other more specialized agents to perform various MDMP tasks. These specialized agents, either based on special roles or Army warfighting functions, are called during various steps and substeps of MDMP, depending on the type of planning task COA-GPT is performing.

To complete the planning cycle, COA-GPT facilitates COA Analysis and Comparison. Once developed, COAs are sent to a simulator to wargame the plans and generate detailed performance metrics. These results are then fed back into COA-GPT, which presents a comparative analysis for the staff and Commander to evaluate. Based on these metrics, the Commander can approve a COA, or the staff can provide feedback to modify the COAs and re-analyze them in the simulator. After a final COA is selected, COA-GPT assists the planning staff in producing the final orders for mission execution, thus supporting the entire MDMP from initial data ingestion to orders production.

At any point during the planning process, the C2 staff can inspect and/or refine the generated MDMP products before approving them, ensuring that decision authority and military judgment remain in human hands. Additionally, a key benefit of our system being driven by LLMs is that the planning staff or Commander can direct or give guidance to COA-GPT, and the system will refine its outputs based on that feedback.

The multi-agent architecture of COA-GPT allows the planning staff to directly interact with multiple specialized LLMs and external tools to execute the MDMP. This agentic design mirrors the structure of a human military planning staff, with specialized staff handling distinct aspects of the MDMP. Beyond its planning capabilities, COA-GPT serves as a comprehensive data management system for all mission and doctrinal documents as it transforms unstructured data and constructs a structured database containing all relevant planning information. COA-GPT also maintains a planning common operating picture (COP) to display its products, manages running estimates of battlefield conditions and resource availability, and generates the final fighting products necessary for mission execution. This integrated approach keeps all planning elements aligned throughout the MDMP, provides Commanders and staff with the most current information for decision-making, and accelerates planning by using the same process the Army already employs on higher echelons.

3.1 Products Generated by COA-GPT

During the **Receipt of Mission** step the operational echelon receives its initial tasking from higher headquarters (HHQ). This includes the higher headquarters' Operations Order (OPORD) with its annexes and accompanying PowerPoint products describing the mission statement (task and purpose), a preliminary timeline, a list of allocated assets or operational constraints, and all other relevant data and intelligence HHQ has compiled. COA-GPT ingests and synthesizes this unstructured information, converting it into a structured format for analysis. It then conducts a preliminary mission assessment to produce the critical data points required for the Initial Warning Order (WARNO) #1. Specifically, COA-GPT processes the task organization, generates the initial timeline, key planning considerations, task to subordinate units, and compiles the Commander's initial guidance. These outputs are then used to formulate and disseminate WARNO #1, officially informing subordinate units and staff and commencing the detailed planning process.

In the **Mission Analysis** step, COA-GPT undertakes a systematic analysis of the Higher Headquarters' (HHQs) orders or plans. This involves an initial Intelligence Preparation of the Operational

Environment (IPOE) to clearly outline operational variables. The process begins with a specialized Geospatial Agent that first analyzes all available map imagery to extract key tactical information, such as terrain features, units on the map, key terrain, mobility corridors, and likely engagement areas. This structured geospatial data is then passed to a Main Agent, which fuses it with information extracted from all available text-based documents and doctrinal publications. With this comprehensive dataset, the Main Agent then executes the formal IPOE process. It starts by describing the environmental effects on operations, then evaluates the threat's capabilities to identify High-Value Targets (HVTs), and culminates in determining and prioritizing the most likely and most dangerous enemy Courses of Action.

Once IPOE is complete, the system identifies and categorizes specified, implied, and essential tasks; thoroughly reviews assets and evaluates resource availability and shortfalls; determines operational constraints; establishes factual data points; formulates necessary assumptions; initiates detailed risk assessment and management processes; and develops initial Commander's Critical Information Requirements (CCIRs) and Essential Elements of Friendly Information (EEFIs). Furthermore, COA-GPT prepares an initial Information Collection (IC) plan, assesses timelines and available time management strategies, and develops both a proposed problem statement and a mission statement. The resulting analysis culminates in issuing Initial Warning Order #2, providing further operational clarity to subordinate commands.

For **Course of Action Development**, COA-GPT continuously refines operational facts and assumptions based on updated situational assessments. It generates diverse and viable COA options, systematically arrays forces according to operational needs, refines each COA's conceptual framework, and assigns relevant headquarters responsibilities. COA-GPT then formulates comprehensive COA statements, and a decision support matrix. For each COA, the system also automatically generates a detailed synchronization matrix. This provides a comprehensive timeline that synchronizes the tasks and purposes of all subordinate units alongside critical decision points and key logistical actions, such as the establishment of command posts and re-supply points. This initial matrix serves as an interactive baseline, allowing users to manually adjust task timings, add new tasks directly in the interface, or request a revised version from COA-GPT based on high-level feedback, as shown in Appendix A.2.

COA sketches are automatically generated as the system progresses through COA development, providing visual representations of all tasks and phases of the proposed COAs for the scenario. These sketches are interactively plotted on the planning Common Operational Picture (COP) — an integrated, dynamic map interface powered by ArcGIS [2]. This map prominently displays units, operational boundaries, strategic objectives, phase lines, and anticipated enemy positions using standardized military symbology, thereby facilitating clear communication and operational alignment.

The planning process continues with **Course of Action Analysis**, where the generated COAs are exported to a simulation environment to be wargamed against anticipated enemy COAs. Based on the resulting performance metrics, COA-GPT then conducts **Course of Action Comparison**. It automatically generates a decision matrix with adjustable weights that contrasts each COA against the evaluation criteria established during Mission Analysis, providing a data-driven recommendation for the staff, as shown in Appendix A.2.

This recommendation directly supports the planning staff in making final refinements before the official **Course of Action Approval**. Once the Commander selects the final COA, the system concludes the MDMP cycle by assisting with **Orders Production**, automatically generating the complete Operation Order (OPORD) to disseminate the plan.

3.2 Human Staff Processes and Interactions

COA-GPT is designed to nest with the iterative and collaborative nature of current military staff interactions while also supporting decentralized collaboration, accelerating the MDMP as a whole. It is hosted as a web-based interface that can be deployed locally or across distributed machines, offering flexibility in various operational environments. Through this interface, users can rapidly progress through the MDMP, generate planning products, and engage with COA-GPT via voice or text, mimicking the informal dialogue typical of traditional planning teams. A sample screenshot

[2] ArcGIS by Esri: `https://www.arcgis.com/index.html`

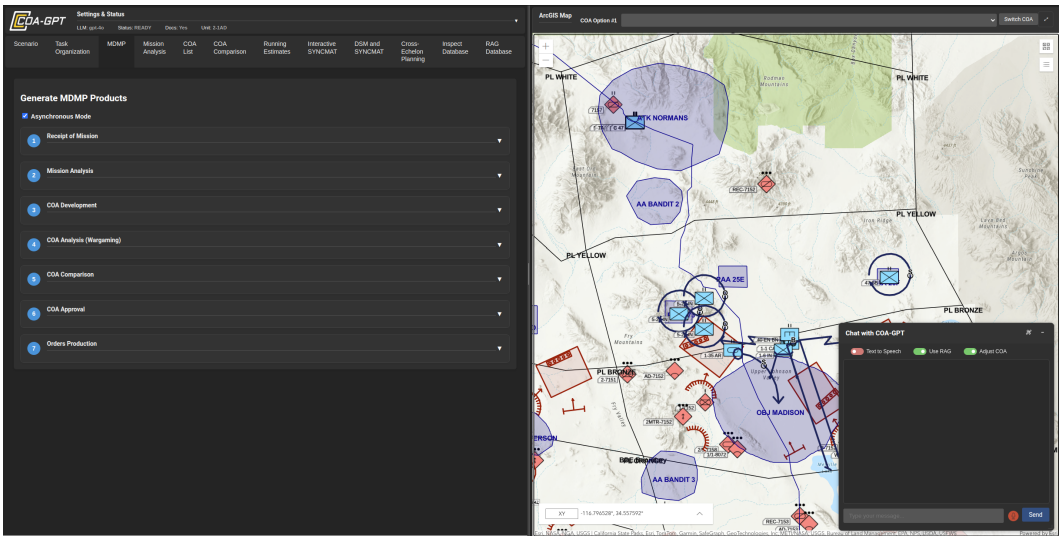

Figure 3: Sample screenshot of COA-GPT's main interface page illustrating the Generation tab, the chat interface, and a sample COA Sketch displayed on the Common Operating Picture (COP).

of COA-GPT's main interface page illustrating the Generation tab, the chat interface, and a sample COA Sketch is shown in Figure 3.

Central to the interface is the Generation tab, structured explicitly around the doctrinal MDMP workflow. Here, staff can initiate product generation such as Warning Orders (WARNOs), COA statements, and synchronization matrices simply by selecting the desired MDMP step and clicking a "Generate" button.

Once a draft product is generated, staff have multiple options to tailor the results. They can directly approve outputs, provide feedback prompting COA-GPT to regenerate refined versions, or manually edit the outputs themselves. Approved products naturally transition COA-GPT to the next MDMP substep, ensuring doctrinal coherence and responsiveness throughout the process.

To bolster trust and transparency, COA-GPT explicitly references source materials such as military doctrine, previous operational products, or specific mission documents, and articulates the reasoning processes behind each decision or recommendation. COA-GPT provides an initial draft to support staff work but is not an authoritative source of truth. Commanders and staff retain full responsibility for reviewing, refining, and approving outputs, ensuring a collaborative human-AI decision-making process.

4 Evaluating COA-GPT

In this evaluation, personnel from the Army's Science & Technology Branch of the Mission Command Capability Development and Integration Directorate (Mission Command CDID) and Army University directly engaged with COA-GPT to plan a mission for the Tropic Tortoise scenario, a military training scenario extensively employed by institutions such as the Command and General Staff College (CGSC) and the Army War College. Participants utilized COA-GPT to generate doctrinal MDMP products and iteratively refined these products based on human feedback. Subsequently, the staff carefully evaluated the final outputs to measure the tool's effectiveness in accelerating Mission Analysis and Course of Action Development during the MDMP.

While COA-GPT addresses the full Military Decision-Making Process (MDMP), this report focuses specifically on products generated during Steps 1 through 3 (Receipt of Mission, Mission Analysis, and COA Development) for the "Tropic Tortoise" scenario, as these steps are the most data-intensive and critical to subsequent planning decisions.

The Tropic Tortoise scenario involves simulated combat operations conducted by the 25th Infantry Division (25 ID) in a contested environment. The operational context centers around offensive

operations within the Mojave Desert, featuring complex multi-domain challenges and enemy capabilities. The primary adversary in this scenario is the 715th Combined Arms Brigade (Heavy Mechanized), representing a conventional, mechanized enemy force with robust defensive tactics and advanced anti-access/area-denial (A2AD) capabilities. Participants assumed the role of planners for the 2nd Brigade, 1st Armored Division (2/1 AD) task organized to the 25 ID, tasked with developing actionable courses of action to seize and fix objectives. Figure 6 in Appendix A.1 illustrates the main area of operation (AO) for this scenario together with the COA Sketch provided by Higher Headquarters, the 25 ID.

Evaluation criteria focused on critical aspects of planning quality, including doctrinal accuracy, completeness, consistency, effectiveness, and overall user satisfaction. Participants graded COA-GPT outputs across multiple MDMP substeps, employing a structured rating scale ranging from 1 ("Not at all accurate or effective") to 5 ("Very accurate, complete, and effective"). More general questions guiding this evaluation included:

1. “How satisfied are users with COA-GPT’s performance in facilitating MDMP? Explain Reasoning",
2. “How much did you trust the information COA-GPT provided?",
3. “How effective is COA-GPT in pulling data verbatim from corpus documents, i.e., Commander’s intent from higher headquarters order? Explain Reasoning",
4. “What level of confidence do human analysts and Commanders have in the outputs provided by COA GPT? Explain Reasoning", and
5. “How effective and useful is COA GPT in supporting/conducting MDMP, compared to a traditional approach? Explain reasoning".

Participants provided qualitative explanations alongside their quantitative assessments, detailing their reasoning and highlighting specific strengths or shortcomings observed during interaction with the tool. Additionally, participants were asked to provide insights into how COA-GPT could influence future Doctrine, Organization, Training, Materiel, Leadership and Education, Personnel, Facilities, and Policy (DOTMLPF-P), further demonstrating its practical implications for future integration into Army operational planning. All human evaluation data was collected using the Behavioral Observations Logging Toolkit (BOLT) [6] platform.

5 Results: Subject Matter Expert Evaluation

Following their use of COA-GPT, participants were surveyed regarding their satisfaction with the system, trust in the information provided, effectiveness of data extraction from source documents, confidence in the outputs, and overall utility compared to traditional MDMP approaches. Results are summarized in Figures 4 and 5 for Mission Analysis and COA Development outputs, respectively. A representative sample of these generated outputs is available for review in Appendix A.2

User satisfaction with COA-GPT’s facilitation of the Military Decision Making Process was above average for Mission Analysis and higher for COA Development. Participants valued the tool’s ability to provide “50 to 70% solutions" that accelerated staff processes and allowed them to focus on refinement rather than starting from scratch. The initial outputs could be easily adjusted to align with Commander guidance, making COA-GPT a valuable starting point for streamlining the early planning stages and reducing staff workload.

Trust in COA-GPT’s information showed similar results, though participants maintained appropriate skepticism. While confidence was generally adequate, trust was undermined by the non-deterministic nature of outputs generated by LLMs, leading to inconsistent trust levels among users who recognized the need for continued human oversight to validate and correct generated information.

COA-GPT demonstrated strong capability in extracting data from corpus documents. Participants particularly valued its ability to extract essential information from base orders, categorize it for quick reference, and summarize higher headquarters’ concepts of operations. However, users noted that the system needed to improve how it synthesizes this information into meaningful operational insights.

When compared to traditional MDMP approaches, COA-GPT proved effective as a productivity and efficiency tool that substantially reduced initial workloads and time commitments for staff officers.

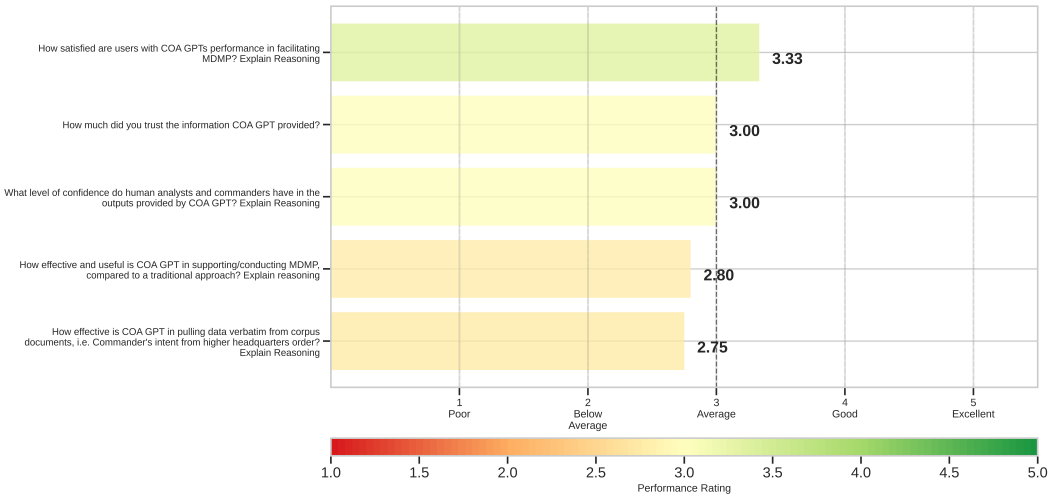

Figure 4: Human evaluation results of COA-GPT when supporting the military staff perform the Mission Analysis step of the MDMP.

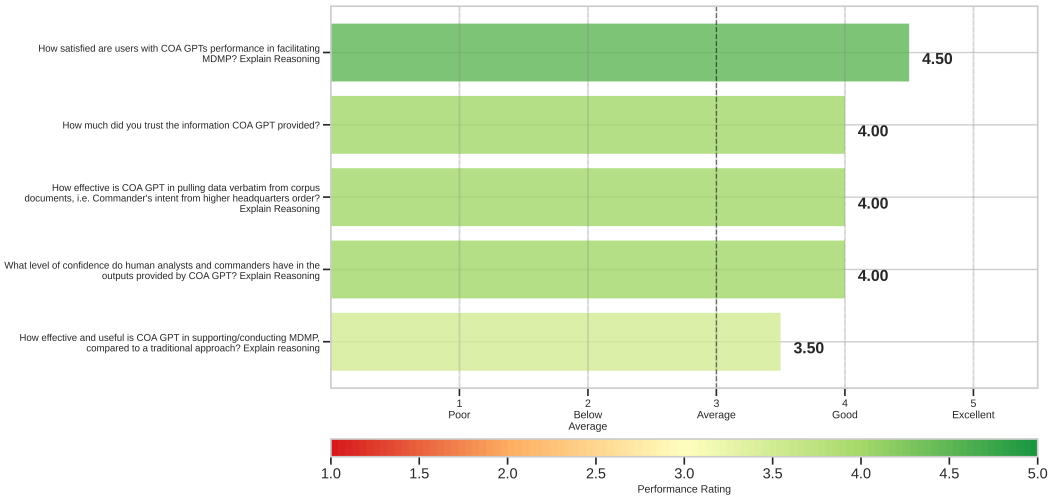

Figure 5: Human evaluation results of COA-GPT when supporting the military staff perform the COA Development step of the MDMP.

Participants valued the tool's speed, adaptability to feedback, and incorporation of doctrine-based reasoning, while acknowledging that human oversight remains necessary to validate, refine, and approve the final MDMP products.

Finally, the evaluation highlighted COA-GPT's capability to identify inconsistencies within planning documents, which emerged as an unexpected benefit for the military staff. The system effectively detected discrepancies within the base scenario corpus, allowing planners to quickly recognize and address contradictions or misalignments between different operational orders and supporting documents that might otherwise require extensive manual cross-referencing. This feature provided immediate practical value to planning teams by automatically flagging potential planning conflicts that could impact mission execution if left unresolved.

6 Discussion

6.1 COA-GPT Performance

COA-GPT directly addresses the emerging battlefield realities outlined earlier, where the deliberate tempo of traditional MDMP processes will likely to be proved insufficient. By leveraging the latest LLM developments, COA-GPT significantly accelerates military planning, transforming what traditionally requires extensive staff resources and prolonged timelines into a rapid, streamlined process. This acceleration enables Commanders and their staffs at higher echelons to exploit fleeting windows of opportunity during a mission, effectively overcoming the mismatch between deliberate MDMP planning and the compressed decision timelines imposed by sophisticated adversaries in contested environments.

Furthermore, by automating the most resource-intensive tasks, COA-GPT supports doctrinal rigor and thorough analysis during the MDMP, enabling planning staff to maintain or enhance their planning effectiveness even in mobile and dispersed command post scenarios.

By providing transparency in decision-making processes, clearly linking generated outputs to doctrinal references and source documents, COA-GPT builds staff confidence and maintains human oversight. This ensures that Commanders retain ultimate control over decision-making, balancing the computational power of AI with the critical judgment of human planners. Ultimately, COA-GPT's integration represents a significant doctrinal evolution, maintaining the agility, responsiveness, and decision advantage essential for successful operations in near-future multi-domain battlefields.

6.2 DOTMLPF-P Implications

The evaluation of COA-GPT revealed significant implications across the DOTMLPF-P framework, highlighting both opportunities and challenges for integrating AI-enabled planning into Army operations.

- **Doctrine**: Participants anticipated MDMP processes becoming more visually oriented, with AI systems generating products directly from graphical inputs. However, they emphasized that planners should first master conventional MDMP before relying on AI tools.
- **Organization**: If systems like COA-GPT consistently produce accurate outputs, personnel requirements could be reduced. At the same time, units would need to place greater emphasis on data literacy across the force.
- **Training**: Effective use of AI requires deliberate training in data literacy and critical thinking, enabling staff to evaluate outputs rather than accept them uncritically.
- **Materiel**: Participants highlighted the need to accelerate AI procurement timelines and ensure secure integration into existing Army networks.
- **Leadership and Education**: Professional Military Education would need adjustments to prepare leaders at all levels to validate AI outputs, recognize limitations, and apply judgment in operational contexts.
- **Personnel**: Staff structures may evolve toward smaller teams composed of more highly skilled personnel with strong data and analytical capabilities.
- **Facilities**: Future facilities would need to connect data systems seamlessly across domains and echelons, while maintaining robust operational security.
- **Policy**: Participants underscored the importance of policies that address risks in AI-generated recommendations for combat decisions, clarify accountability, and set guidelines for secure cross-domain data access.

7 Conclusions

By enabling the accelerated mission analysis, development, and refinement of COAs and the MDMP as a whole, COA-GPT can reshape the relationship between Commanders, staffs, and the planning processes. With COA-GPT, higher echelons can maintain decision advantage even in contested environments with DDIL communications, enabling more rapid and informed command decisions that

seize fleeting opportunities and respond dynamically to battlefield developments. At the operational level, integrating COA-GPT into C2 processes significantly accelerates planning by automating key steps of the MDMP. This increased efficiency allows command teams to generate and refine courses of action faster, supporting more agile and responsive decision-making.

Going beyond the planning process, if integrated with operations' execution, COA-GPT can continuously update and refine plans as the operational situation evolves. Unlike traditional planning cycles, our system can potentially ingest real-time battlefield information, identify planning discrepancies, and rapidly generate branches and sequels that capitalize on emergent opportunities or mitigate unexpected threats. This creates a near-continuous planning cycle that merges seamlessly with execution, potentially reshaping how military operations are conceptualized, planned, and conducted across the full spectrum of conflict. Importantly, COA-GPT remains a staff support tool: commanders and planners retain full responsibility for evaluating and approving all outputs. By augmenting rather than replacing human judgment, COA-GPT represents a balanced approach to integrating artificial intelligence into the inherently human domain of military decision-making.

Acknowledgements

The authors thank the U.S. Mission Command Battle Lab for their guidance on MDMP instruction, for steering COA-GPT's development toward Army operational requirements, for supporting COA-GPT's assessment, and for the insightful discussions regarding the integration of AI planning tools into future Army Command and Control concepts. We are also grateful to the entire BOLT team at the U.S. Army DEVCOM Army Research Laboratory for providing the data collection platform that supported the military subject matter expert evaluation effort.

Funding

Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-23-2-0072. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

References

[1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

[2] Jonathan Boron and Chris Darken. Developing combat behavior through reinforcement learning in wargames and simulations. In *2020 IEEE Conference on Games (CoG)*, pages 728–731. IEEE, 2020.

[3] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. Sparks of artificial general intelligence: Early experiments with gpt-4. *arXiv preprint arXiv:2303.12712*, 2023.

[4] Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025.

[5] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv e-prints*, pages arXiv–2407, 2024.

[6] CJ Garneau, BE Hoffman, and NE Buchler. Behavioral Observations Logging Toolkit (BOLT): Initial deployed prototypes and usability evaluations. ccdc data & analysis center–devcom reports. aberdeen proving ground md, usa, 2020.

[7] Vinicius G Goecks and Nicholas Waytowich. 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)*, pages 01–10. IEEE, 2024.

[8] Vinicius G Goecks, Nicholas Waytowich, Derrik E Asher, Song Jun Park, Mark Mittrick, John Richardson, Manuel Vindiola, Anne Logie, Mark Dennison, Theron Trout, et al. On games and simulators as a platform for development of artificial intelligence for command and control. *The Journal of Defense Modeling and Simulation*, 20(4):495–508, 2023.

[9] Larry Ground, Alexander Kott, and Ray Budd. A knowledge-based tool for planning of military operations: The coalition perspective. In *Proceedings of the Second International Conference on Knowledge Systems for Coalition Operations*, pages 195–203, 2002.

[10] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

[11] Headquarters, Department of the Army. *Planning and Orders Production*. United States Army, 2024.

[12] Alexander Kott, Larry Ground, Ray Budd, Lakshmi Rebbapragada, and John Langston. Toward practical knowledge-based tools for battle planning and scheduling. In *AAAI/IAAI*, pages 894–899, 2002.

[13] Zhong-Zhi Li, Duzhen Zhang, Ming-Liang Zhang, Jiaxin Zhang, Zengyan Liu, Yuxuan Yao, Haotian Xu, Junhao Zheng, Pei-Jie Wang, Xiuyi Chen, et al. From system 1 to system 2: A survey of reasoning large language models. *arXiv preprint arXiv:2502.17419*, 2025.

[14] Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*, 2024.

[15] P Narayanan, M Vindiola, S Park, A Logie, N Waytowich, M Mittrick, J Richardson, D Asher, and A Kott. First-Year Report of ARL Director's Strategic Initiative (FY20-23): Artificial Intelligence (AI) for Command and Control (C2) of Multi Domain Operations. Technical Report ARL-TR-9192, Adelphi Laboratory Center (MD): DEVCOM Army Research Laboratory (US), May 2021.

[16] Daniel C Ruiz and John Sell. Fine-tuning and evaluating open-source large language models for the army domain. *arXiv preprint arXiv:2410.20297*, 2024.

[17] Peter J Schwartz, Daniel V O'Neill, Meghan E Bentz, Adam Brown, Brian S Doyle, Olivia C Liepa, Robert Lawrence, and Richard D Hull. AI-enabled wargaming in the military decision making process. In *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications II*, volume 11413, page 114130H. International Society for Optics and Photonics, 2020.

[18] John R Surdu and Kevin Kittka. Deep green: Commander's tool for coa's concept. In *Computing, Communications and Control Technologies*, volume 34, pages 45–51, 2008.

[19] Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024.

[20] Nicholas Waytowich, Sean L Barton, Vernon Lawhern, Ethan Stump, and Garrett Warnell. Grounding natural language commands to StarCraft II game states for narration-guided reinforcement learning. In *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications*, volume 11006, page 110060S. International Society for Optics and Photonics, 2019.

[21] Nicholas Waytowich, James Hare, Vinicius G Goecks, Mark Mittrick, John Richardson, Anjon Basak, and Derrik E Asher. Learning to guide multiple heterogeneous actors from a single human demonstration via automatic curriculum learning in starcraft ii. In *Artificial intelligence and machine learning for multi-domain operations applications IV*, volume 12113, pages 283–293. SPIE, 2022.

A APPENDIX

A.1 SCENARIO INFORMATION: 25ID COA SKETCH

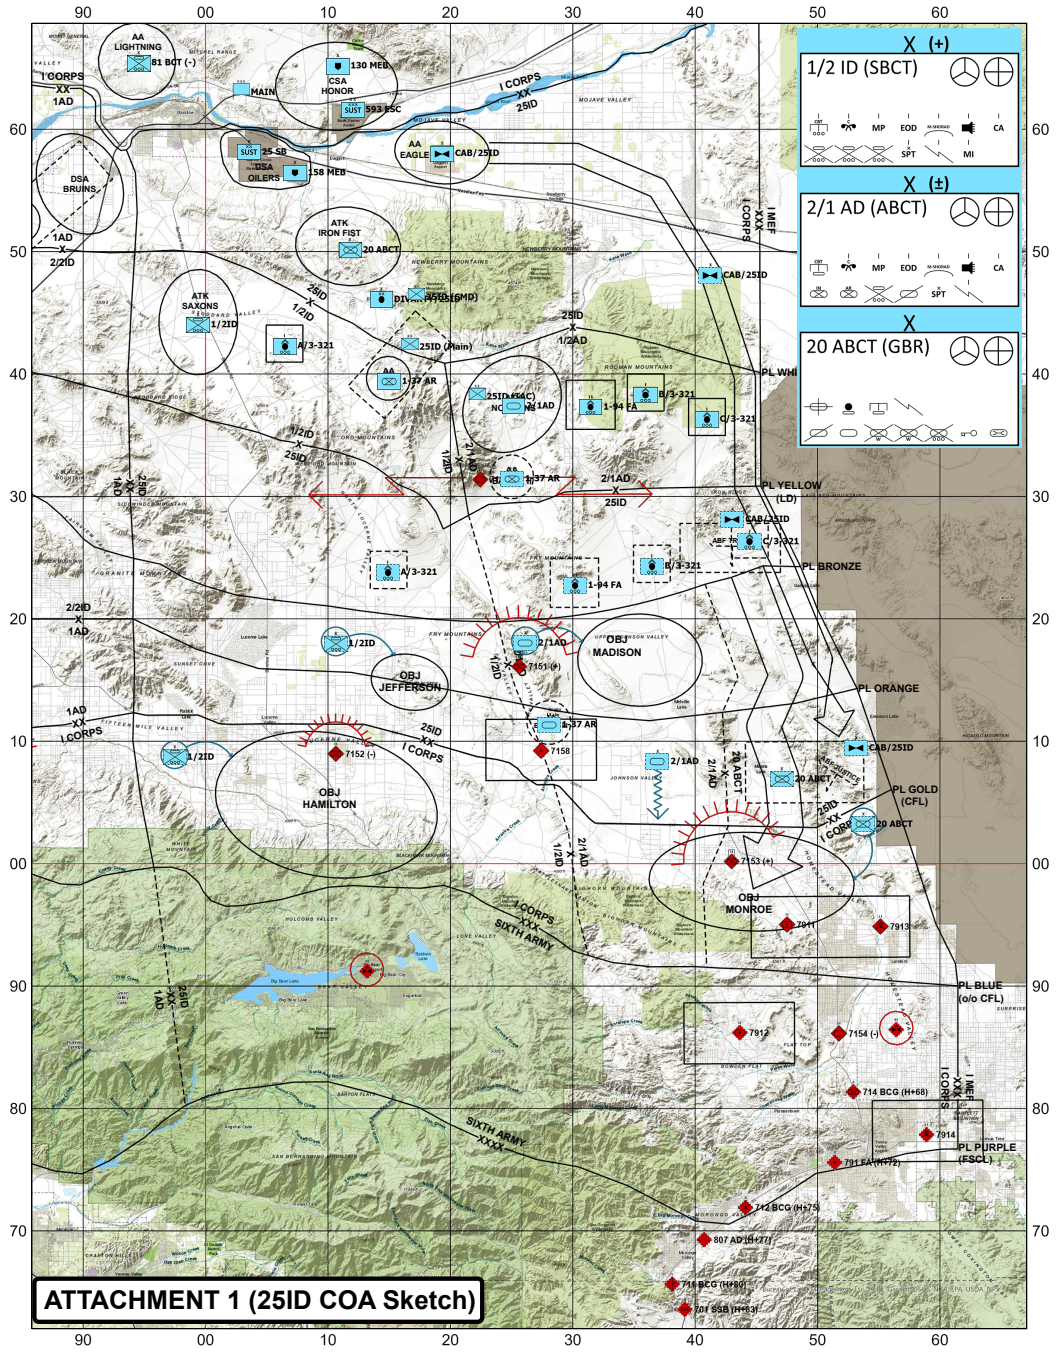

Figure 6: Area of Operation (AO) and Higher Headquarters' (HHQ), the 25th Infantry Division (25 ID), COA Sketch for the Tropic Tortoise training scenario used to evaluate COA-GPT's role in supporting the planning staff of the 2nd Brigade, 1st Armored Division (2/1 AD) executing the MDMP.

A.2 SAMPLE COA-GPT PRODUCTS

The following sections showcase a curated selection of products to demonstrate the capabilities of COA-GPT within the Tropic Tortoise scenario. The products shown are illustrative examples and do not represent the full range of outputs. All products were generated by COA-GPT alone, with no human supervision or correction. Furthermore, they were generated across multiple, distinct planning sessions and are not from a single, continuous event.

The samples are organized into two subsections corresponding to their MDMP phase. First, text-based products from Mission Analysis are displayed, followed by a graphical product from COA Development.

Text-Based Products (Mission Analysis):

- Analysis of Higher HQ Order;
- Commander's Critical Information Requirements (CCIRs);
- Mission and Problem Statements;
- Specified, Implied, and Essential Tasks;
- Warning Order #1 (Figures 7 to 14).

Graphical Product (COA Development):

- Updated Task Organization (Figure 15);
- Synchronization Matrix (Figure 16);
- COA Comparison Table (Figure 17).

A.2.1 Select Text Products: Mission Analysis

Analysis of Higher HQ Order

HHQ COMMANDER'S INTENT

Detailed Source
Details: I Corps COA Commander's intent para; 25 ID OPORD para 3-a (Commander's Intent)
Filename: I Corps COA Statement.docx; 25ID_OPORD_2029-25.pdf

Implications For Unit
2-1 AD must seize OBJ MADISON quickly, destroy enemy reconnaissance north of PL BRONZE, and fix enemy depth forces on OBJ MONROE to enable the division decisive operation while protecting corps LOCs.

One Level Up
25 ID's purpose is to prevent the enemy from gaining access to the I-40 and I-15 corridors. Key tasks are (1) identify 715 BCG frontier battle positions to degrade the continuity of their defense, (2) prevent 715 BCG reserves from reinforcing their depth zone, and (3) forward-position sustainment nodes to maintain momentum. End-state: 715 BCG withdrawn south of OBJ HAMILTON and OBJ MONROE; 25 ID defends PL BLUE and consolidates gains.

Two Levels Up
I Corps' purpose is to prevent the Olvanan-Southern Military Theater from seizing ground lines of communication into LAS VEGAS. Key tasks are (1) degrade enemy air-defense and fire-strike groups along the SAN BERNARDINO MOUNTAINS, (2) prevent 16 ACG from disrupting friendly sustainment near BARSTOW, (3) reposition corps sustainment nodes forward to support three brigades in HOMESTEAD VALLEY, and (4) prevent enemy reconnaissance in the corps rear. End-state: 71 ACG unable to defend OBJ MONROE, 16 ACG contained south of PL WHITE, I Corps postured defensively north of the passes and ready to transition to stability operations.

HHQ CONCEPT OF OPERATIONS

Detailed Source
Details: OPORD para 3b Concept of Operations (close, deep, rear areas; scheme of maneuver and sustainment)
Filename: 25ID_OPORD_2029-25.pdf

Implications For Unit
2-1 AD must achieve rapid penetration, facilitate forward passage of 20 ABCT, and then fix enemy depth forces to enable the decisive envelopment.

Main Effort
Initially 2-1 AD (seize OBJ MADISON); main effort transitions to 20 ABCT after OBJ MADISON is secured.

Overview
25 ID conducts a frontal attack south from PL YELLOW, with 2-1 AD and 1-2 ID attacking abreast through enemy frontier defense positions, followed by a shift of main effort to 20 ABCT (GBR) to envelop and seize OBJ MONROE; division then transitions to a defense on PL BLUE.

Supporting Efforts
1-2 ID clears OBJ JEFFERSON then OBJ HAMILTON to secure division west flank; CAB/25 ID provides aviation attack and MEDEVAC; DIVARTY conducts counter-fire and SEAD.

Sustaining Operations
25 SB pushes sustainment from DSA OILERS, later jumps to AA BANDIT II; 158 MEB clears MSRs/ASRs and establishes detainee and decon sites.

HHQ MISSION ANALYSIS

Assigned Tasks

- Degrade enemy air-defense and fire-strike groups (I Corps to 25 ID)
- Prevent 16 ACG disruption of BARSTOW sustainment (I Corps to 25 ID)
- Seize OBJ MADISON (25 ID to 2-1 AD)
- Fix enemy on OBJ MONROE (25 ID to 2-1 AD)

Detailed Source
Details: OPORD para 2 Mission; para 3b Scheme of Maneuver (assignment of main effort)
Filename: 25ID_OPORD_2029-25.pdf

Implications For Unit
2-1 AD's assigned tasks directly support the corps and division purpose by removing enemy control of OBJ MADISON and preventing reinforcement of OBJ MONROE.

One Level Up
Mission – NLT H+96, 25 ID seizes key terrain vicinity LUCERNE VALLEY and JOHNSON VALLEY to deny 71 ACG the ability to posture for follow-on operations toward LAS VEGAS.

Purpose
Deny the enemy the ability to project combat power north of the SAN BERNARDINO MOUNTAINS and secure friendly LOCs into the Mojave Desert.

Two Levels Up
Mission – NLT H+96, I Corps seizes HOMESTEAD VALLEY to prevent O-SMT access to the I-40/I-15 corridors.

Commander's Critical Information Requirements (CCIRs)

PIR-1: Will the 715 BCG defend OBJ MADISON with more than one CA-Bn, and where are the main defensive positions and obstacles along AA2a/AA2b and AA3a/AA3b? (DP1a, DP2)

PIR-2: Has the 7154 CA-Bn (-) or 852 AV attack helicopters initiated movement to conduct a counterattack against 2-1AD ABCT's main effort or support elements? (DP1c)

PIR-3: What is the current location and firing status of enemy SR-4 and PCL-09 artillery batteries (7158 FA, 7911 FA, 7913 MRL) capable of massing fires on 2-1AD ABCT's breach and support areas? (DP2, DP3a)

PIR-4: What is the status and location of enemy air defense radars (Type-306B/HT-233) and HQ-17A batteries within range of 2-1AD ABCT's planned aviation operations? (DP3a, DP3b)

FFIR-1: Has 2-1AD ABCT crossed PL BRONZE and initiated the breach on AA2a/AA2b? (DP2)

FFIR-2: Has 2-1AD ABCT seized OBJ MADISON and established a foothold south of PL GOLD? (DP3a)

FFIR-3: Is 2-1AD ABCT effectively fixing enemy forces on OBJ MONROE from the north to enable 20 ABCT's envelopment? (DP3b)

FFIR-4: Is 2-1AD ABCT's combat power below 60

FFIR-5: What are the planned times and routes for 2-1AD ABCT's sustainment convoys and reserve force commitment?

Mission and Problem Statements

Problem Statement: How does 2-1AD ABCT rapidly breach and seize OBJ MADISON, fix the 715 BCG's main defensive group on OBJ MONROE, and transition to defense along PL BLUE, in compartmentalized desert terrain dominated by key passes and ridgelines, against a layered mechanized defense with massed fires, obstacles, and local counterattack, given limited reconnaissance, extended sustainment lines, and the need to maintain combat power above 60%, within a 96-hour window, while minimizing civilian disruption and coordinating with non-English-speaking populations and limited local governance?

Mission Statement: 2nd Brigade, 1st Armored Division (2-1AD) Armored Brigade Combat Team will seize OBJ MADISON and fix enemy forces on OBJ MONROE from the north NLT H+48 in the assigned area of operations in order to enable the division's decisive operation by denying the 715 BCG the ability to defend north of the SAN BERNARDINO MOUNTAINS.

Specified, Implied, and Essential Tasks

Specified Tasks

- Seize OBJ MADISON to divert enemy forces away from AXIS HORNET.
- Fix enemy forces on OBJ MONROE from the north to enable 20 ABCT envelopment.
- Clear PAAs 25E, 25F, and 173 to protect 3-321 FA and 1-94 FA.

Implied Tasks

- Conduct reconnaissance to identify enemy positions and adjust plans.
- Coordinate with adjacent units to ensure synchronized operations.
- Establish supply lines to support sustained operations.

Essential Tasks

- Seize OBJ MADISON to divert enemy forces and support the main effort.

- Fix enemy forces on OBJ MONROE to enable 20 ABCT envelopment.
- Clear PAAs 25E, 25F, and 173 to protect artillery assets.

[UNCLASSIFIED]

Copy 1 of 10 copies

Change from verbal orders: None

Issuing headquarters: 2nd Brigade, 1st Armored Division (2-1AD ABCT)

Place of issue: AA BANDIT (11SNU 1494 3966)

Date-time group of signature: 1200Z D+13 2029

WARNING ORDER #1 (U)

(U) References: OPORD 12029-01 (MOJAVE MANEUVERS) – I Corps; 25ID OPORD 2029-25 (TROPIC TORTOISE); Map Sheet: Mojave Desert, 1:50,000

(U) Time Zone Used Throughout the OPLAN or OPORD: TANGO (UTC-7)

(U) Task Organization: See Annex A (Task Organization)

1. (U) Situation.

a. (U) Area of Interest. The area of interest includes the SAN BERNARDINO MOUNTAINS, YUCCA VALLEY, JOHNSON VALLEY, LUCERNE VALLEY, and all enemy assembly areas and likely avenues of approach into the 2-1AD ABCT area of operations.

b. (U) Area of Operations. 2-1AD ABCT operates in the central sector of the 25ID AO, oriented south, with the FEBA along the ORD MOUNTAINS (11SNU 1441 3593) and extending south toward PL BLUE. The AO is compartmentalized by mountain ranges and includes key terrain at LUCERNE VALLEY and JOHNSON VALLEY.

c. (U) Enemy Forces. The enemy 715 Combined Arms Brigade (Heavy Mechanized) is conducting a hasty defense in the vicinity of YUCCA VALLEY and JOHNSON VALLEY, reinforced with artillery, engineers, and air defense. Expect layered defenses: Frontal Blocking Zone (PL WHITE–PL BRONZE), Frontier Defense Zone (PL BRONZE–PL GOLD), Depth Defense Zone (PL GOLD–PL BLUE), and Rear Defense Zone (PL BLUE–PL PURPLE). Enemy reserves and counterattack forces are likely in depth, with significant artillery and air defense threats. Enemy is expected to employ scatterable mines and complex obstacles.

d. (U) Friendly Forces. 25ID is the higher headquarters, with the mission to seize key terrain in LUCERNE VALLEY and JOHNSON VALLEY to deny the 71 ACG the ability to posture for follow-on operations toward LAS VEGAS. Adjacent units: 1/2ID (SBCT) to the west, 20 ABCT (GBR) to the east, and CAB/25ID in support. I Corps' purpose is to prevent O-SMT from seizing ground lines of communication into LAS VEGAS.

e. (U) Interagency, Intergovernmental, and Nongovernmental Organizations. Provincial government and NGOs are supporting humanitarian assistance in unoccupied areas. IDP camps are

Figure 7: Page 1/8 of Sample WARNORD #1 generated by COA-GPT for the Tropic Tortoise scenario (no human feedback or review).

located at NEWBERRY SPRINGS and HINKLEY. Civil authorities maintain limited control north of PL BRONZE.

f. (U) Civil Considerations. The AO includes small urban centers (LUCERNE VALLEY, YUCCA VALLEY) with significant civilian populations. 60% of the population has evacuated; 40% remain. Avoidance of combat in population centers is a priority. Infrastructure and supply lines may be affected by displaced civilians and humanitarian operations.

2. (U) Mission. 2-1AD ABCT attacks to seize OBJ MADISON and fix enemy forces on OBJ MONROE NLT H+96 in order to enable 25ID to seize key terrain in LUCERNE VALLEY and JOHNSON VALLEY and deny the 71 ACG the ability to posture for follow-on operations toward LAS VEGAS.

3. (U) Execution.

a. (U) Initial Commander's Intent. The purpose of this operation is to rapidly seize OBJ MADISON, fix enemy forces on OBJ MONROE, and maintain momentum to enable the division's decisive operation. Key tasks: conduct reconnaissance and counter-reconnaissance, breach enemy obstacles, destroy enemy mechanized and armored forces, and protect the brigade's flanks. End state: 2-1AD ABCT controls OBJ MADISON, fixes enemy on OBJ MONROE, and is postured to defend along PL BLUE.

b. (U) Concept of Operations. To be determined. 2-1AD ABCT will conduct reconnaissance and counter-reconnaissance, attack to seize OBJ MADISON, and prepare to fix enemy forces on OBJ MONROE in support of the division's main effort.

c. (U) Tasks to Subordinate Units.

- 1-1 CAV SQ (organic): Screen along PL YELLOW to PL BRONZE in order to provide early warning and destroy enemy reconnaissance forces to enable freedom of maneuver for the brigade.

- 1-6 IN (organic): Seize assigned sector of OBJ MADISON in order to destroy enemy mechanized forces and enable the brigade's penetration of the Frontier Defense Zone.

- 1-35 AR (organic): Seize assigned sector of OBJ MADISON in order to destroy enemy armored forces and enable the brigade's penetration of the Frontier Defense Zone.

- 5-20 IN (attached): Seize assigned sector of OBJ MADISON in order to destroy enemy mechanized forces and enable the brigade's penetration of the Frontier Defense Zone. Maintain close coordination with adjacent British armor to prevent fratricide.

- 4-27 FA (SP) (organic): Attack by fire enemy artillery and defensive positions in order to suppress, neutralize, or destroy enemy forces and enable maneuver of brigade elements.

- 40 EN BN (organic): Breach enemy obstacles along the axis of advance in order to enable rapid maneuver of brigade combat power onto OBJ MADISON.

- 47 SB (organic): Sustain brigade operations in order to enable continuous combat operations and maintain combat power.

Figure 8: Page 2/8 of Sample WARNORD #1 generated by COA-GPT for the Tropic Tortoise scenario (no human feedback or review).

- HHC 2/1 (organic): Command and control brigade operations in order to synchronize subordinate unit actions and maintain operational tempo.

- 369 CM CO (HR) (OPCON): Conduct CBRN reconnaissance and decontamination in order to enable freedom of maneuver and protect brigade forces from CBRN threats.

- 856 MP CO (GS) (OPCON): Guard main supply routes and conduct detainee operations in order to protect sustainment operations and enable freedom of movement.

- 756 OD CO (EOD) (OPCON): Clear explosive hazards in order to enable safe movement and operations for brigade forces.

- A/6-56 ADA (M-SHORAD) (OPCON): Protect brigade forces from enemy air threats in order to enable freedom of maneuver and preserve combat power.

- 303 PSYOPS CO (OPCON): Influence civilian and enemy actions in order to support brigade operations and minimize interference with military objectives.

- A/411 CA BN (OPCON): Support interaction with local authorities and civilian population in order to facilitate civil-military operations and maintain freedom of action.

d. (U) Coordinating Instructions.

- OPORD to follow NLT 2000 D+14.

- H-Hour and LD to be published in OPORD.

- Units begin movement preparations, LOGSTAT reporting, and PCC/PCI.

- Initial planning meeting at 2-1AD TOC, 1200 D+14.

- Maintain OPSEC; no movement south of PL YELLOW without brigade approval.

- All vehicles will display IR marker panels for fratricide avoidance.

- MOPP2 south of PL YELLOW; review CBRN procedures.

- Submit requests for Class IV/V and maintenance support to 47 SB.

- Report all changes in readiness or significant events immediately to brigade TOC.

- Rules of Engagement: Consistent with Khorathidin government’s declaration of martial law; non-combatants can be detained and transported by coalition forces for up to 24 hours, then released or turned over to Khorathidin police.

- Risk Reduction: Develop lateral routes to move contaminated personnel and equipment to ASR SUNRISE.

- Fire Support Coordination Measures: FSCL is PL PURPLE; I Corps CFL is PL BLUE (O/O PL PURPLE); 25ID CFL is PL GOLD (O/O PL BLUE).

Figure 9: Page 3/8 of Sample WARNORD #1 generated by COA-GPT for the Tropic Tortoise scenario (no human feedback or review).

- Air and Missile Defense: Warning Status YELLOW; Weapon Control Status TIGHT.

- Personnel Recovery: Units are responsible for personnel recovery in their AO.

4. (U) Sustainment.

- Priority of supply: Class III (fuel), V (ammunition), VIII (medical), IX (repair parts).

- Brigade Support Battalion (47 SB) will maintain 3 days of supply on hand; coordinate with 25 SB for resupply.

- Units will depart TAA with full combat loads and 15 DOS of CLIII-P.

- Submit LOGSTAT and maintenance status NLT 1800 daily.

- Casualty evacuation is unit responsibility to brigade AXPs; Role II/III at division level.

- Prepare to establish ATHP and FARP locations; report to brigade S4.

- Water: Bulk water is the primary source; no bottled water will be distributed outside CSA COURAGE.

- JSLIST: One per Soldier; OCIE issues restricted to emergency requirements.

- Submit requests for CLIV to 25 SB SPO NLT 72 hours prior to required delivery.

5. (U) Command and Signal.

- No change to current command relationships.

- Succession of command: 2-1AD CDR, then 1-6 IN CDR, then 1-35 AR CDR.

- Brigade TOC at AA BANDIT (11SNU 1494 3966) until further notice.

- Maintain current comms plan; updates to be published in OPORD.

- Provide LNOs to adjacent units as directed.

Annexes:

Annex A – Task Organization

Annex B – Intelligence

Annex C – Operations

Annex D – Fires

Annex E – Protection

Annex F – Sustainment

Figure 10: Page 4/8 of Sample WARNORD #1 generated by COA-GPT for the Tropic Tortoise scenario (no human feedback or review).

Annex H – Signal

Annex K – Civil Affairs

Annex L – Information Collection

Distribution:

CDR, 1-1 CAV SQ

CDR, 1-6 IN

CDR, 1-35 AR

CDR, 5-20 IN

CDR, 4-27 FA (SP)

CDR, 40 EN BN

CDR, 47 SB

CDR, HHC 2/1

CDR, 369 CM CO (HR)

CDR, 856 MP CO (GS)

CDR, 756 OD CO (EOD)

CDR, A/6-56 ADA (M-SHORAD)

CDR, 303 PSYOPS CO

CDR, A/411 CA BN

S1, S2, S3, S4, S6

25ID TOC

[UNCLASSIFIED]

Annexes Definition:

Annex A – Task Organization

Headquarters Unit: 2/1 AD (Brigade HQ)

- HHC 2/1 (Armor), Organic, Organic

- 369 CM CO (HR), OPCON, Chemical

Figure 11: Page 5/8 of Sample WARNORD #1 generated by COA-GPT for the Tropic Tortoise scenario (no human feedback or review).

- 856 MP CO (GS), OPCON, Military Police
- 756 OD CO (EOD), OPCON, Ordnance
- A/6-56 ADA (M-SHORAD), OPCON, Air Defense
- 303 PSYOPS CO, OPCON, Psychological Operations
- A/411 CA BN, OPCON, Civil Affairs
- 1-6 IN (Infantry Battalion), Organic, Organic
- HQ 1-6 (Company), Organic, Infantry
- A 1-6 (Company), Organic, Infantry
- B 1-6 (Company), Organic, Infantry
- C 1-6 (Company), Organic, Infantry
- FWD SPT G 47 (Company), Organic, Support
- 1-35 AR (Armor Battalion), Organic, Organic
- HQ 1-35 (Company), Organic, Armor
- A 1-35 (Company), Organic, Armor
- B 1-35 (Company), Organic, Armor
- C 1-35 (Company), Organic, Armor
- FWD SPT H 47 (Company), Organic, Support
- 5-20 IN (Mechanized Infantry Battalion), Attached, Attached
- HQ 5-20 (Company), Attached, Mechanized Infantry
- A 5-20 (Company), Attached, Mechanized Infantry
- B 5-20 (Company), Attached, Mechanized Infantry
- C 5-20 (Company), Attached, Mechanized Infantry
- FWD SPT H 296 (Company), Attached, Support
- 1-1 CAV SQ (Reconnaissance Squadron), Organic, Organic
- HQ 1-1 (Troop), Organic, Reconnaissance
- A 1-1 (Troop), Organic, Reconnaissance

Figure 12: Page 6/8 of Sample WARNORD #1 generated by COA-GPT for the Tropic Tortoise scenario (no human feedback or review).

- B 1-1 (Troop), Organic, Reconnaissance
- C 1-1 (Troop), Organic, Reconnaissance
- D 1-1 (Troop), Organic, Reconnaissance
- FWD SPT D 47 (Company), Organic, Support
- 4-27 FA (SP) (Field Artillery Battalion), Organic, Organic
- HQ 4-27 (Battery), Organic, Field Artillery
- A 4-27 (Battery), Organic, Field Artillery
- B 4-27 (Battery), Organic, Field Artillery
- C 4-27 (Battery), Organic, Field Artillery
- TA/HQ 4-27 (Battery), Organic, Target Acquisition
- FWD SPT F 47 (Company), Organic, Support
- 40 EN BN (Engineer Battalion), Organic, Organic
- HQ 40 (Company), Organic, Engineer
- A 40 (Company), Organic, Engineer
- B 40 (Company), Organic, Engineer
- 510 EN CO (CEC-A), OPCON, Engineer
- C 40 (Company), Organic, Engineer
- D 40 (Company), Organic, Engineer
- FWD SPT E 47 (Company), Organic, Support
- 47 SB (Support Battalion), Organic, Organic
- HQ 47 (Company), Organic, Support
- A 47 (Company), Organic, Support
- B FLD 47 (Company), Organic, Maintenance
- C 47 (Company), Organic, Support

Annex B – Intelligence

Figure 13: Page 7/8 of Sample WARNORD #1 generated by COA-GPT for the Tropic Tortoise scenario (no human feedback or review).

[Placeholder: To be published in OPORD]

Annex C – Operations

[Placeholder: To be published in OPORD]

Annex D – Fires

[Placeholder: To be published in OPORD]

Annex E – Protection

[Placeholder: To be published in OPORD]

Annex F – Sustainment

[Placeholder: To be published in OPORD]

Annex H – Signal

[Placeholder: To be published in OPORD]

Annex K – Civil Affairs

[Placeholder: To be published in OPORD]

Annex L – Information Collection

[Placeholder: To be published in OPORD]

[UNCLASSIFIED]

Figure 14: Page 8/8 of Sample WARNORD #1 generated by COA-GPT for the Tropic Tortoise scenario (no human feedback or review).

A.2.2 Select Graphical Products: COA Development

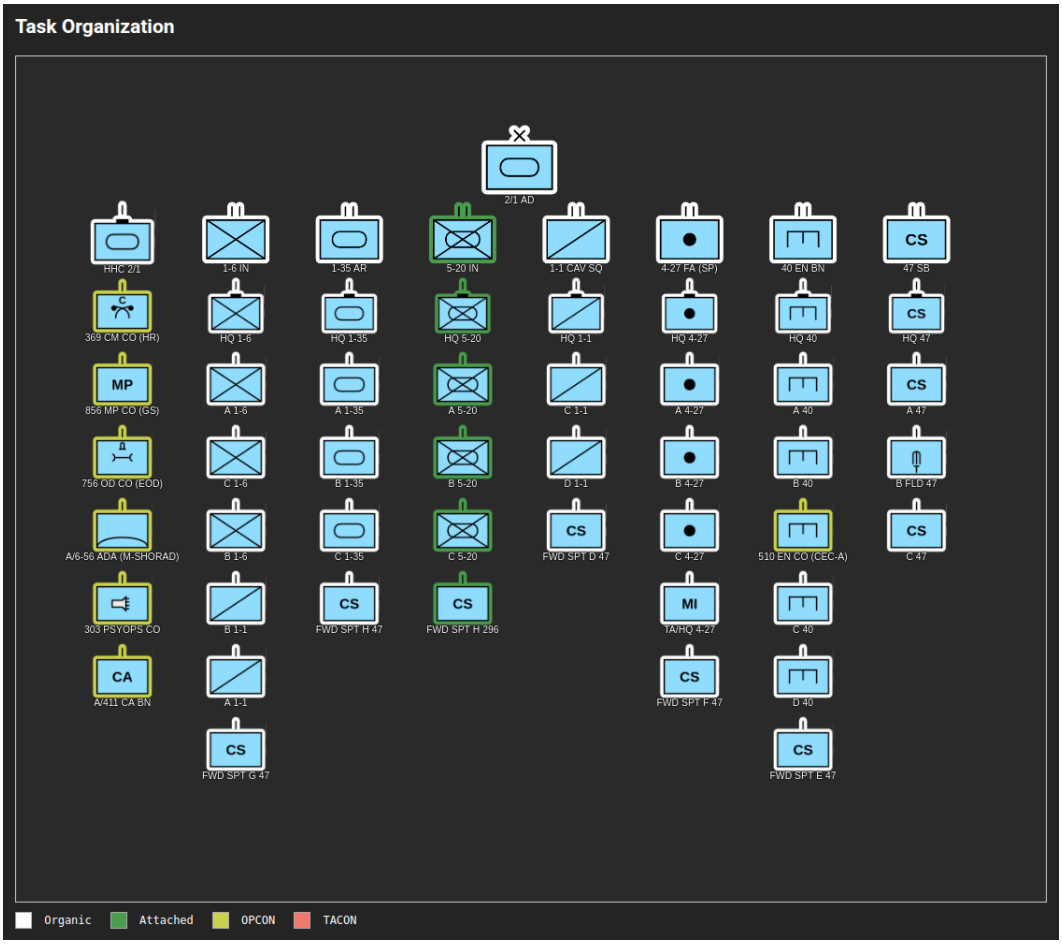

Figure 15: Sample updated Task Organization for one of the COAs generated by COA-GPT for the Tropic Tortoise scenario with proposed reinforcements to the Main Effort (no human feedback or review).

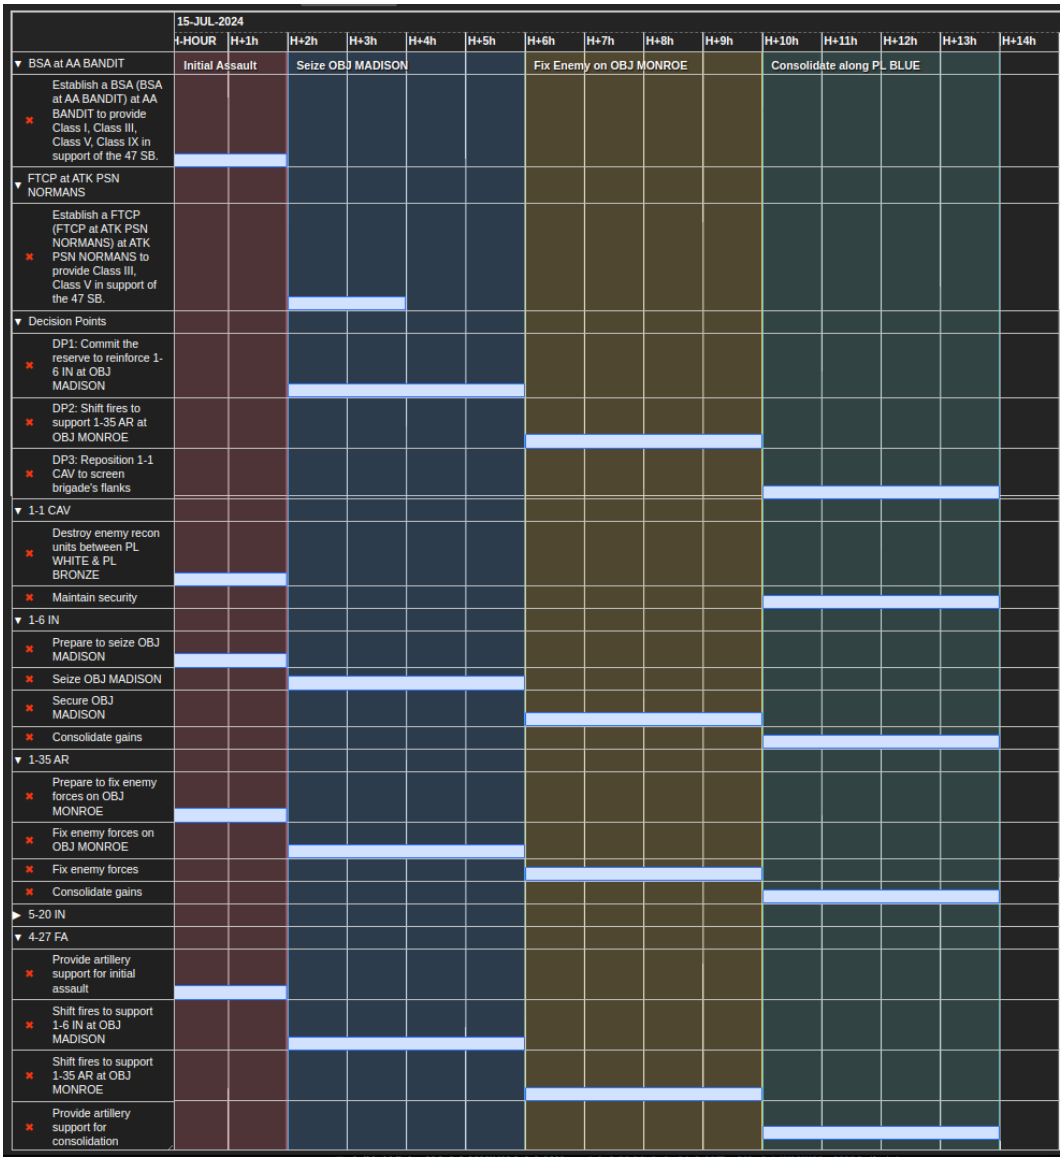

Figure 16: Sample Synchronization Matrix for one of the COAs generated by COA-GPT for the Tropic Tortoise scenario (no human feedback or review).

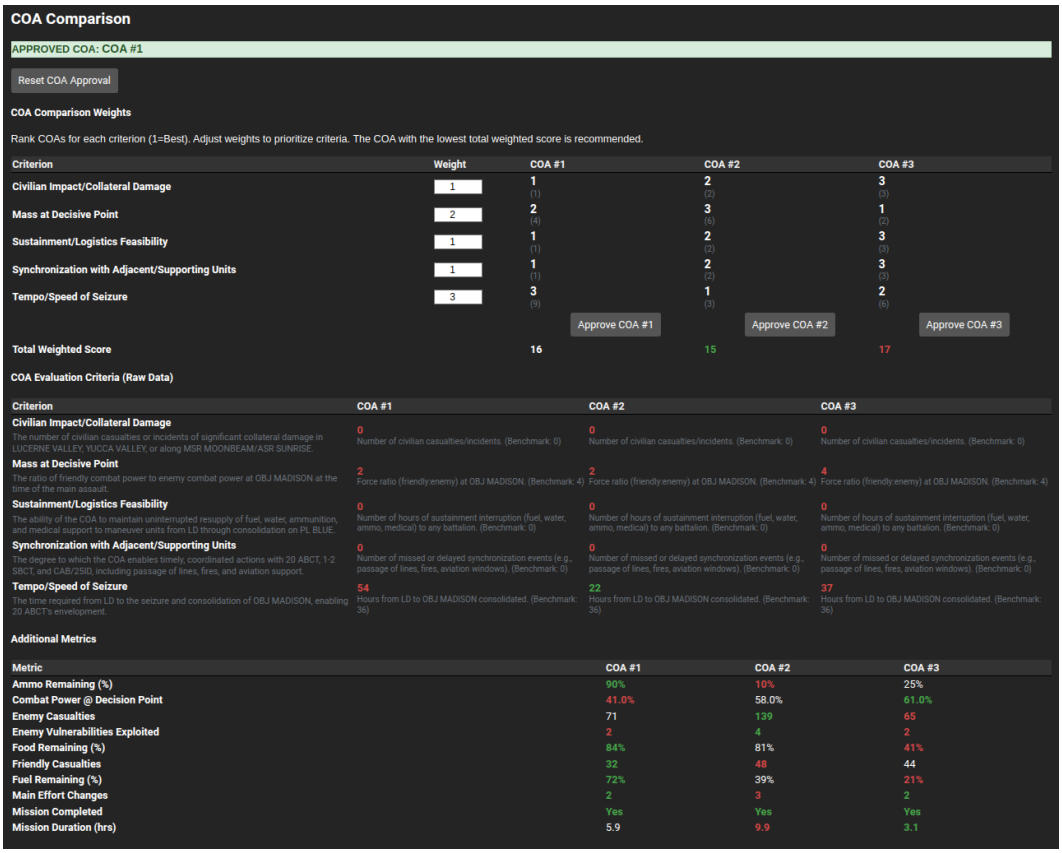

COA Comparison

APPROVED COA: COA #1

Reset COA Approval

COA Comparison Weights

Rank COAs for each criterion (1=Best). Adjust weights to prioritize criteria. The COA with the lowest total weighted score is recommended.

Criterion	Weight	COA #1	COA #2	COA #3
Civilian Impact/Collateral Damage	1	1 (1)	2 (2)	3 (3)
Mass at Decisive Point	2	2 (4)	3 (6)	1 (2)
Sustainment/Logistics Feasibility	1	1 (1)	2 (2)	3 (3)
Synchronization with Adjacent/Supporting Units	1	1 (1)	2 (2)	3 (3)
Tempo/Speed of Seizure	3	3 (9)	1 (3)	2 (6)
		Approve COA #1	Approve COA #2	Approve COA #3
Total Weighted Score		16	15	17

COA Evaluation Criteria (Raw Data)

Criterion	COA #1	COA #2	COA #3
Civilian Impact/Collateral Damage The number of civilian casualties or incidents of significant collateral damage in LUCERNE VALLEY, YUCCA VALLEY, or along MSR MOONBEAM/ASR SUNRISE.	0 Number of civilian casualties/incidents. (Benchmark: 0)	0 Number of civilian casualties/incidents. (Benchmark: 0)	0 Number of civilian casualties/incidents. (Benchmark: 0)
Mass at Decisive Point The ratio of friendly combat power to enemy combat power at OBJ MADISON at the time of the main assault.	2 Force ratio (friendly:enemy) at OBJ MADISON. (Benchmark: 4)	2 Force ratio (friendly:enemy) at OBJ MADISON. (Benchmark: 4)	4 Force ratio (friendly:enemy) at OBJ MADISON. (Benchmark: 4)
Sustainment/Logistics Feasibility The ability of the COA to maintain uninterrupted resupply of fuel, water, ammunition, and medical support to maneuver units from LD through consolidation on PL BLUE.	0 Number of hours of sustainment interruption (fuel, water, ammo, medical) to any battalion. (Benchmark: 0)	0 Number of hours of sustainment interruption (fuel, water, ammo, medical) to any battalion. (Benchmark: 0)	0 Number of hours of sustainment interruption (fuel, water, ammo, medical) to any battalion. (Benchmark: 0)
Synchronization with Adjacent/Supporting Units The degree to which the COA enables timely, coordinated actions with 20 ABCT, 1-2 SBCT, and CAB/2SID, including passage of lines, fires, and aviation support.	0 Number of missed or delayed synchronization events (e.g., passage of lines, fires, aviation windows). (Benchmark: 0)	0 Number of missed or delayed synchronization events (e.g., passage of lines, fires, aviation windows). (Benchmark: 0)	0 Number of missed or delayed synchronization events (e.g., passage of lines, fires, aviation windows). (Benchmark: 0)
Tempo/Speed of Seizure The time required from LD to the seizure and consolidation of OBJ MADISON, enabling 20 ABCT's envelopment.	54 Hours from LD to OBJ MADISON consolidated. (Benchmark: 36)	22 Hours from LD to OBJ MADISON consolidated. (Benchmark: 36)	37 Hours from LD to OBJ MADISON consolidated. (Benchmark: 36)

Additional Metrics

Metric	COA #1	COA #2	COA #3
Ammo Remaining (%)	90%	10%	25%
Combat Power @ Decision Point	41.0%	58.0%	61.0%
Enemy Casualties	71	139	65
Enemy Vulnerabilities Exploited	2	4	2
Food Remaining (%)	84%	81%	41%
Friendly Casualties	32	48	44
Fuel Remaining (%)	72%	39%	21%
Main Effort Changes	2	3	2
Mission Completed	Yes	Yes	Yes
Mission Duration (hrs)	5.9	9.9	3.1

Figure 17: Sample COA Comparison Table for the COAs generated by COA-GPT for the Tropic Tortoise scenario (no human feedback or review).