

HR0011SB20254-07
Improving Battle Planning through AI
Frequently Asked Questions (FAQs)

1. Are there particular types of scenarios or environments that you're most interested in exploring during this effort?

A: The government is interested in the ability to develop compositional models reflective of the underlying dynamic systems that can be composed for multiple scenarios.

2. Could you share a bit about any specific challenges or pain points you've observed with current COA adjudication methods?

A: The primary pain points are speed and accuracy for scaling purposes. Using reduced order models the government seeks to accelerate adjudication for horizontal and vertical scaling to many times faster than real time.

3. Is there a preference or requirement regarding integration and composability with existing simulation or war-gaming frameworks?

A: Responsive solutions will generalize for composition in future simulation and war-gaming frameworks over standard APIs.

4. The majority of the opportunity is tailored to specialized COA/ROM software. Will there be an opening for an architecture opportunity?

A: Responsive solutions will be general and adaptable to many next-generation government COA generation and adjudication frameworks.

5. Could DARPA clarify whether the envisioned ROM adjudication architecture is expected to include adversarial modeling or co-evolution of blue/red COAs, or should the focus remain on evaluating blue COAs only?

A: The adjudication architecture should be capable of modeling both red, blue, and green COAs evolving at the same time in a ROM representation of a multi-physics simulation environment.

6. Does the government have a preference for how data heterogeneity is resolved within federated models (e.g., ontology alignment, schema harmonization, or data fusion heuristics), or is this left to the proposer?

A: This is left to the proposer.

7. Are LLMs, knowledge graphs, or neural-symbolic hybrid systems considered responsive under the topic's intent, provided they are used to support composable ROM-based adjudication?

A: Provided the methods accelerate to many times beyond real time, and can be validated against trusted physics simulator, proposers may submit any technology solution they believe meets criteria. It should be noted, accuracy is a key metric, however.

8. Would integration of reinforcement learning (e.g., for adversarial red COA generation or COA refinement under uncertainty) be aligned with the objectives of this topic?
A: See Question 7.
9. Is there DARPA interest in incorporating models for psychological, civil, or information-based effects—such as morale, civilian sentiment, or influence operations—into the ROM adjudication process?
A: Yes, DIMEFIL modeling are responsive.
10. Can you confirm whether the government envisions real-time ROM-based adjudication in denied, degraded, intermittent, or low-bandwidth (DDIL) environments, or if centralized/cloud environments are assumed?
A: The government envisions the ROM-based adjudication in next-generation war gaming environments.
11. Does DARPA expect direct interoperability with any existing planning systems (e.g., GCCS, JOPES), or is demonstrating API-driven composability sufficient?
A: Demonstrating a generalizable API is more desirable than interoperability with a single planning system.
12. How will the accuracy of COA adjudication be measured during Phase II? Are baseline datasets or scoring references available?
A: Against the original data sets or simulation environments used to generate the ROMs.
13. Are there specific expectations for modeling strategic effects across the DIMEFIL (Diplomatic, Information, Military, Economic, Financial, Intelligence, Law Enforcement) spectrum?
A: Expectations are primarily with respect to model fidelity.
14. Are there potential integration or transition pathways envisioned with programs such as SCEPTER, or expectations for working within classified, coalition, or interagency environments?
A: Multiple transition pathways with DARPA and broader DoD needs are possible for a successful effort.
15. What echelon of command is envisioned as the primary user of the COA engine (e.g., tactical unit commander, joint task force, theater-level)? The COAs generated would differ significantly depending on strategic vs. tactical planning intent.
A: Responsive proposals will address strategic, operational, and tactical levels.
16. Is the target battlespace primarily Army, Air Force, Navy—or joint/combined with coalition forces? We've observed significant differences when planning with ROK or other allied partners during SOF wargaming.
A: Responsive proposals will address joint, all-domain, warfare. Partners and allies modeling is also desired.

17. Will models explicitly account for variables such as weather, terrain, electromagnetic spectrum conditions, and time-of-day? These are key in real-world COA feasibility but often missing from abstract models.

A: Responsive proposals will address variables represented by the core modeling and sim environments being accelerated by ROM development.

18. How does the government envision modeling uncertainty or incomplete information? Is there an expectation to incorporate probabilistic fog-of-war mechanisms into ROM-based adjudication?

A: ROM-based adjudication systems should faithfully represent all characteristics of their original environments, including sensors, uncertainty, and incomplete information.

19. Will the government permit adaptive tolerance ranges in MOR algorithms that refine fidelity dynamically—without full retraining or manual recalibration—to enable COA adjudication at speeds five orders of magnitude faster than real time?

A: Yes.

20. Can the government clarify whether there is a preferred modeling environment or simulation platform (e.g., AFSIM, WARSIM, JADC2 testbeds) with which the ROM architecture should interoperate in Phase II demonstrations?

A: Proposed methods should be general enough to apply to any modeling and simulation environment.

21. To what extent should the proposed ROMs be able to ingest both real-time streaming inputs and pre-staged simulation data? Is there a performance or fidelity tradeoff that should be optimized for?

A: The proposed ROMs should be able to function as an accelerated version of the original environment for adjudication purposes, operating faster than real time, and with bounded error.

22. Does DARPA envision the composition operator and ROMs being reusable across operational domains (air, land, cyber), or is tailoring expected per domain due to differing physics models and planning logics?

A: The ROMs should be compositional across operational domains.

23. Are there specific standards or protocols DARPA recommends for model composability or federated model integration (e.g., HLA, FMI, DDS, or other simulation standards)?

A: Proposers should ensure the generality of their APIs for interfacing.

24. For DIMEFIL representation in federated planning, is there any weighting or prioritization across the dimensions that the adjudication process should emphasize (e.g., M vs. I vs. F)?

A: Responsive proposals will validate their solutions across multiple aspects of the DIMEFIL to show generalizability.

25. Will the government provide sample COA datasets or adversarial simulation logs to calibrate initial ROM performance during early prototyping?

A: If requested, COA datasets can be provided for existing M&S environments.

26. Is it acceptable for the Phase II prototype to rely on CPU-only inference pipelines for portability and containerization, or are GPU-accelerated ROM pipelines encouraged where latency is critical?

A: Provided the necessary acceleration is achieved with minimal error, no hardware restrictions exist.