



WAR ON THE ROCKS

**When the world's at stake,
go beyond the headlines.**

National security. For insiders. By insiders.

Join War on the Rocks and gain access to content trusted by policymakers, military leaders, and strategic thinkers worldwide.

BECOME A MEMBER



COGS OF WAR

AI Is Being Misunderstood as a Breakthrough in Planning. It's Not.



CHRISTOPHER DENZEL

February 26, 2026

In the age of AI, the scarcest resource in headquarters is no longer time. It is, rather, the willingness to say no.

Artificial intelligence is moving rapidly into military planning staffs because it compresses routine cognitive labor. AI excels at absorbing guidance, reorganizing complex material, and producing clear strategic language at speed. This feels like a qualitative advance, creating the impression that

planning itself has become easier. But this impression misleads. The risk of AI-enabled planning is that it will produce plausible constructs that obscure where judgment is required, creating the illusion that analytic completeness can substitute for prioritization.

AI is seen as “raising the floor” by making it easier to produce adequate products. That is true. Yet AI also “collapses the median” by increasing the relative cost of real insight. As AI-enabled planning begin to inform real-world operations, the temptation is to treat complete answers as sufficient, without interrogating whether they represent the right answers to the hard questions of what to resource, what to defer, and what risk to accept.

This temptation confuses exhaustiveness with clarity. At the campaign level, these characteristics find themselves in tension. Failure rarely stems from insufficient information but instead from the failure to impose priority where competing objectives cannot all be satisfied. AI synthesizes brilliantly, but synthesis cannot resolve this tension.

The following observations reflect lessons from sustained use of large language models by myself and other planners at U.S. Forces Japan. During transformation campaign planning, AI was used not as an experiment but rather embedded directly into workflows for synthesis, reframing, drafting, and iteration. Used inside live planning cycles, the limits – and advantages – of AI emerged repeatedly.

AI changes campaign planning, but in a narrower way than often assumed. It raised the floor, collapsing the time and effort required to generate and revise internally coherent constructs. However, the same features that give AI its greatest potential in military planning also allow unresolved choices to hide behind orderly structure — collapsing the median.

AI's value lies less in producing a strong theory of the campaign than in rapidly revealing which theories do not survive contact with competing demands and constraints. In practice, AI should be used to present commanders not with complete solutions but with multiple internally-consistent framings, with planners exposing where those frames break — forcing decisions about priority, risk, and subordination.

Operational Art Is Bounded by Judgment, Not Information

Operational art is governed by a constraint that no amount of analytic power can remove because the most consequential decisions in campaign design are made where causality is uncertain. Priority is not a question of quality or quantity of data — it is a question of judgment.

In a campaign, the problem is not understanding what should be done, but deciding what matters more when objectives compete, authorities overlap, resources are finite, and adversaries adapt. These judgments are consequential, yet cannot be proven correct in advance. A campaign design can only be judged plausible given assumptions that may or may not hold. Operational art, therefore, relies on causal reasoning under uncertainty, not on predictive models that promise a “correct answer.”

This is why additional information and more sophisticated synthesis do not reliably reduce uncertainty in such environments. As complexity grows, the marginal value of detail declines and the importance of focusing effort increases. Additional analysis cannot resolve which risks are acceptable, which objectives must prevail, or which failures can be absorbed. That judgment — the imposition of priority where demands cannot all be satisfied — is an inherent, and human, function of command.

This constraint applies regardless of the tool set in use. Campaign design has always employed analytic aids — intelligence preparation, modeling, red-teaming — but none displaced the requirement for accountable choice. Before AI, these processes were labor-intensive and iterative, limiting how quickly staffs could generate a coherent plan and forcing planners to dwell on complex issues. New technologies may accelerate synthesis and iteration, but they encounter a hard ceiling. Algorithms cannot decide what deserves priority and the speed they promise threatens to crowd out the time required for judgment. Treating those decisions as analytic outputs obscures where judgment must be exercised.

The Temptation of Treating AI as a “Planning Optimizer”

Organizations may be tempted to manage uncertainty through optimization — arranging objectives, tasks, and effects into plans that appear complete, balanced, and institutionally defensible. While this impulse predates AI, what changes is how easily this can now be indulged.

But optimization fails because it solves the wrong problem. At the campaign level, coherence and balance cannot substitute for a causal theory and exhaustiveness cannot replace judgment.

The optimizing mindset produces familiar artifacts: evenly weighted objectives, symmetrical lines of effort, comprehensive task lists, and balanced effects that appear to cover the problem space. AI excels at generating planning constructs that are structurally-sound and defensible. But that coherence is precisely the danger. By smoothing differences and averaging the focus of effort, optimization suppresses priority.

AI makes balanced, defensible plans cheap and plentiful — allowing unresolved priorities to hide behind competent structure. Optimization substitutes internal order for choice and symmetry for judgment. Perfunctory acknowledgment that “judgment cannot be automated” functions less as a constraint than as a permission structure. I’ve seen staffs proceed as if coherent AI-generated recommendations were an acceptable stand-in for decision. Absent observable consequences, that acknowledgement becomes performative rather than constraining. Staffs ought to manifest behavior changes in an auditable way, presenting AI-generated framings alongside analysis of where those framings fail, presenting commanders with not just a proposed course, but highlighting the limits and tradeoffs that require commander’s judgment.

An optimized plan places internal logic over external reality, allowing balance to substitute for focus. The result diffuses responsibility.

This failure is structural, not technical. It cannot be fixed with better data, tighter prompts, or additional iteration. Optimization assumes that importance is latent in the data and can be surfaced through better analysis. Campaign design operates under a different logic. Priority is not discovered in the data, but asserted by an authority. Algorithms cannot determine which objective should dominate, which lines of effort should be subordinate, or which risks are acceptable. Answers to those questions cannot be found in the data—they are a choice based on a commander's judgment and accountability.

When AI Reveals the Limits of Structure

The dynamics described above become most visible where campaigns resist a single, coherent frame. This was evident during campaign planning at U.S. Forces Japan, where the mission set does not align cleanly with assumptions embedded in traditional operational design. U.S. Forces Japan operates at the intersection of military readiness, alliance management, policy execution, and diplomacy — often simultaneously. These responsibilities draw from different authorities, pursue different logics, and answer to different measures of success. They are related, but not contiguous.

Traditional planning struggles in such environments because it presumes that unity can be imposed through structure. Doctrinal constructs (objectives, lines of operation, decisive points, etc.) work best when a mission can be reduced to a dominant logic. In non-traditional missions, that reduction misleads. The problem is a mismatch between the framework being imposed and the underlying reality being described. AI will cleanly apply doctrine to a complex problem, smoothing out where the reality defies doctrinal categorization. When a machine does this, commanders and staffs may never be aware that a square peg has been algorithmically rendered “round” to fit the shape of the hole.

AI rendered this mismatch explicit in campaign planning at U.S. Forces Japan. There, the planning team largely worked on classified networks, using Maven Smart Systems, which provided Anthropic's Claude Sonnet model. Planners also experimented with Ask Sage (which provided access to a variety of models, including those of OpenAI).

When applied to U.S. Forces Japan's campaign design, large language models consistently produced analytically defensible but conceptually fragile frames defining the command's roles, missions, and functions. Each frame cleanly organized tasks, clustered objectives, and generated plausible narratives, yet each failed in a different place, usually by subordinating one role while privileging another. No single construct could accommodate all roles simultaneously without distortion.

This was not a failure of prompting or data. Rather, it was a signal that the mission itself resisted singular explanation. These failures made the underlying reality unmistakable. The command's mission was not only non-traditional, it was also non-contiguous — operating simultaneously at the

policy, strategic, and tactical levels, each governed by distinct logics and supporting relationships.

That breakdown proved productive. Because reframing was cheap, planners could generate competing constructs, observe where each collapsed, and interrogate why. AI did not reveal a “correct” frame, but rather revealed the limits of each available one. Instead of planners providing the “best fit” frame or presenting leading frames as options to select from, planners recommended that the commander accept the structural inconsistency of three roles, each with distinct missions and functions. Paradoxically, by rejecting a single framing of the command’s purpose, the commander preserved concurrent, independent prioritization within each core role. This tension, in fact, clarified that U.S. Forces Japan’s function is not to resolve competing roles into a unified hierarchy, but to integrate a set of irreducible responsibilities into a single headquarters capable of providing unity of effort where unity of command is impossible.

This is where AI “raises the ceiling.” By rapidly producing competent framings, AI strips away false differentiation between plans that only appear distinct. What remains — where each construct breaks or distorts — is precisely where command judgment should be exercised. In that sense, AI raises the ceiling of planning not by deciding, but by making the need for command judgment easier to locate.

This dynamic is not unique to U.S. Forces Japan. It emerges wherever campaign responsibilities resist reduction to a single governing logic. Once optimization is set aside, AI’s practical effect in campaign design becomes clearer and more limited. AI does not resolve uncertainty or displace judgment. It collapses the cost of generating plausible ways of thinking about the problem. The danger comes when the standard for planning is lowered to plausibility—now available in an instant. A more responsible approach to AI-enabled planning instead adopts the harder method of using multiple plausible plans *diagnostically* to surface the tensions that spark real insight.

That acceleration alters the rhythm of planning. Instead of investing disproportionate effort in producing a single polished construct, staffs can explore multiple framings, discard weak ones early, and revisit assumptions more often. But this same speed creates risk. Coherent plans can appear “good enough” to brief before priority has been imposed. Absent deliberate intervention, accelerated tempo favors structure over decision, allowing unresolved choices to persist beneath orderly, balanced frameworks.

Campaigns do not succeed through balance but through deliberate imbalance — privileging some objectives, subordinating others, and accepting specific risks in service of a dominant purpose. Those judgments are not latent in data. They are imposed through a commander’s intent. Seen this way, AI’s symmetry bias becomes diagnostically useful. It exposes exactly where command judgment should intervene. AI can organize the problem and sharpen debate, but it cannot rank what should prevail. Imposing asymmetry on a resistant problem remains inseparable from command itself.

Judgment Requires Sacrifice

The most dangerous failures in AI-enabled planning will not look like technical breakdowns. Plans will remain coherent, language disciplined, and assumptions defensible — providing a comfortable structure that invites consensus, allowing responsibility to diffuse.

Once asymmetry is made explicit, familiar planning comforts become incompatible with how campaign design actually works. These are not transitional frictions or governance gaps, but structural tradeoffs that cannot be avoided.

The first tradeoff is exhaustiveness. Campaigns cannot be both exhaustive and adaptable. When reframing is cheap, the instinct is to keep adding tasks, effects, and lines of effort until the plan feels complete. And large language models will infinitely oblige requests to go deeper. But exhaustive constructs harden assumptions. In contrast, campaign plans should be expected to change.

The second tradeoff is symmetry. Balanced frameworks feel fair, but they defer the decisive act of choosing what should dominate when resources, time, or political capital tighten. Asymmetry requires subordination. Some objectives advance at the expense of others, and some risks are accepted explicitly rather than smoothed away. That imbalance cannot be achieved through refinement. It requires a visible decision to elevate one purpose and accept the degradation of another.

The third tradeoff is plausible deniability. Balanced and exhaustive plans mask hard choices about risks and priorities within structure. When everything is prioritized, nothing is. Imposed asymmetry does the opposite — it makes judgment legible. Declaring a dominant objective, an acceptable risk, a main effort, or a subordinate line of operations ties outcomes to decisions rather than process.

These losses may be mislabeled as process problems, but no refinement of governance, weighting schemes, or review cycles can substitute for the act of establishing clear priorities. Attempts to compensate in this way miss the point. What is being lost is not procedural comfort, but the ability to hide choice behind structure. The discomfort that follows is not evidence of failure. It is evidence that authority is being exercised.

If AI-enabled planning fails, it will not be because the technology was insufficient but because structure became cheap and easy, obscuring where tradeoffs should be made. When AI-generated coherence is allowed to stand in for decision, plans appear disciplined while the hard choices of command recede.

AI supercharges linear planning doctrine. If commanders and staffs stop here, judgment will be obscured behind an increasingly competent structure. But paradoxically, the acceleration AI unlocks in planning should cause commanders and staffs to slow down. Using AI to locate where judgment and priority is most needed takes time.

Christopher Denzel is a Marine Corps strategic and operational planner currently serving in the Strategic Plans Branch (J-56) at U.S. Forces Japan. His work focuses on strategy, campaigning, and command and control modernization in the Indo-Pacific, with extensive experience integrating planning, intelligence, and wargaming in joint and allied environments.

He is a graduate of the U.S. Army School of Advanced Military Studies and the National Intelligence University, and has held operational planning billets across the joint force, including at U.S. Cyber Command and III Marine Expeditionary Force.

The views expressed are his own and do not represent the views or positions of U.S. Forces Japan, the U.S. Marine Corps, the Department of Defense, or any part of the U.S. government.

***Please note, as a matter of house style, War on the Rocks will not use a different name for the U.S. Department of Defense until and unless the name is changed by statute by the U.S. Congress.*

Image: Megan Hearst via [DVIDS](#).



Overview

About

Account

Advertising

Contact

Events and Sponsorships

People

Submissions

Non-Members

Applied History

Battle Studies

Book Reviews

Cogs of War

War by Other Ledgers

War On The Rocks

What's Going On In Iran?

What's Going On In Venezuela?

Members

Membership

Get More War On The Rocks

Support Our Mission And Get Exclusive Content

BECOME A MEMBER

Subscribe to our newsletter

Enter your email address

SUBSCRIBE

By signing up you agree to our data policy

[Privacy Policy](#)

[Terms & Conditions](#)

[Sitemap](#)

Metamorphic Media.