Impossibility Engineering: Teaching AI to Recognize When Adversaries Change the Game

Darrell L. Young[1], Mitchell A. Thornton[1], Jason Teske[2], and James D. Moreland, Jr.[3]

[1]Darwin Deason Institute for Cybersecurity, Southern Methodist University, Dallas, TX 75205, USA
[2]George Washington University, Washington, DC 20052, USA
[3]Virginia Tech, Arlington, VA 22203, USA

Abstract

Multi-domain operations face challenges where optimization fails not because the math is wrong, but because the problem is incorrectly framed. Adversaries exploit this: cyberattackers use social engineering when firewalls prove impenetrable; doctrine inversions create surprise when opponents refuse to fight the expected war; negotiations succeed by reframing positions as interests. This paper introduces Impossibility Engineering—a methodology for detecting when you or your organization may be stuck in the wrong frame, and for discovering reframings that others have used to solve analogous problems. We implement a Context Continuum architecture where a manager LLM coordinates specialized components: one generates radical reframing hypotheses, another critiques them, others retrieve historical analogies from a curated corpus of breakthrough cases. Systems-of-Systems integration reveals constraint inversion opportunities invisible when analyzing components in isolation. We demonstrate applications in negotiations, cyber defense, and tactical planning, offering not a solution but an invitation: the methodology and tools described here are designed to help practitioners recognize and solve their own impossible problems.

Keywords: Impossibility Engineering, frame-lock detection, adversarial reasoning, ChatPack, Context Engineering, large language models, multi-domain operations, reframing methodology

1 Introduction: The Stone That Shouldn't Skip

Consider a specification that seems impossible to satisfy:

> *Levitate a dense object across a horizontal plane for a sustained distance without thrust, lift surfaces, power systems, or active control.*

An engineer reviewing this specification would likely reject it. Dense objects fall. Sustained horizontal travel requires propulsion or aerodynamic lift. The laws of physics appear to forbid the requested behavior.

Yet a child at a lake demonstrates the solution routinely: a stone, skipping across water, travels dozens of meters horizontally while the thrower watches. The stone doesn't violate physics—it exploits a regime (high-velocity impact with a fluid surface at shallow angles) that the specification's framing failed to consider.

This paper is about that gap—the space between "impossible" and "not yet reframed."

1.1 The Defense Relevance

Translate the stone-skipping scenario to multi-domain operations. An intelligence, surveillance, and reconnaissance (ISR) system is configured to detect airborne threats. Its sensors monitor for thrust signatures, radar cross-sections consistent with aircraft, and flight dynamics within expected parameters. An adversary, understanding this configuration, employs a different approach: a dense object, thrown at shallow angle, skips across the sensor's detection zone. The object exhibits none of the signatures the system was designed to detect. The sensors report: no contacts.

The adversary didn't defeat the sensor. They reframed the problem—operating in a regime the sensor's designers hadn't anticipated.

This pattern recurs across domains. Cyberattackers facing hardened network perimeters shift to social engineering.[**?**] Negotiators deadlocked on positions discover movement when they reframe to interests.[**?**] Tacticians facing superior forces refuse the expected engagement, choosing instead to change the nature of the conflict entirely.

1.2 What This Paper Offers

The methodology presented here—Impossibility Engineering—does not claim to solve impossible problems. Rather, it offers practitioners a systematic approach to:

1. **Detect** when a problem may be incorrectly framed (frame-lock detection)
2. **Retrieve** historical cases where analogous reframings succeeded
3. **Generate** candidate reframing hypotheses for human evaluation
4. **Critique** those hypotheses against domain constraints

The goal is not to replace human judgment but to augment it—to surface possibilities that time pressure, cognitive load, or organizational assumptions might otherwise obscure.

We describe the methodology, an architecture for implementing it, and three application domains. We conclude with an invitation: the tools and corpus described here are designed to help you recognize and solve your own impossible problems.

2 The Frame-Lock Problem

When optimization algorithms report that no feasible solution exists, the natural response is to seek better algorithms, more computational resources, or relaxed constraints. But there is a prior question: *Is the optimization even formulated correctly?*

2.1 Combinatorial Explosion as Symptom

Multi-domain operations planning exhibits characteristic complexity. A mission involving eight coalition partners, each with distinct capabilities, rules of engagement, and communication protocols, generates a combinatorial space that overwhelms direct optimization. Planners facing such complexity often report that the problem is "impossible"—meaning intractable within available time and resources.

But combinatorial explosion is sometimes a symptom, not a cause. The explosion may indicate that the planning framework has failed to identify which decisions actually matter.

Our earlier work on multi-agent coordination demonstrated this pattern.[?] A routing problem with m waypoints, optimized jointly across all agents, exhibits $O(m!)$ complexity. But decomposing the problem into strategic decision waypoints (where meaningful choices occur) and tactical motor waypoints (interpolated between decisions) reduced complexity to $O(m^3)$. The "impossible" optimization became tractable—not through better algorithms, but through recognizing that most waypoints weren't decisions at all.

2.2 Adversary Exploitation of Frame-Lock

Adversaries actively seek frame-lock in opposing systems. History's most decisive victories resulted not from superior forces but from attackers who recognized constraints their opponents assumed were real. Table 1 summarizes landmark cases.

Table 1: Historical examples of Impossibility Engineering achieving strategic surprise

Year	**Operation**	**"Impossible" Constraint**	**Reframing**
218 BC	Hannibal's Alps	Alps impassable with elephants	Undefended route > fortress wall
1940	Ardennes	Forest impassable for armor	Impassable = undefended
1941	Pearl Harbor	Harbor too shallow for torpedoes	Engineer wooden stabilizers
1950	Inchon	Tides/seawalls prevent landing	Impossible = unexpected
1976	Entebbe	Hostages 2,500 miles away	Distance = no expectation

The pattern is consistent: strategic surprise emerges not from superior capability but from operating outside the defender's frame. MacArthur captured it precisely regarding Inchon: "The very arguments you have made as to the impracticabilities involved will tend to ensure for me the element of surprise."[?]

2.3 Deception, Surprise, and Intelligence Analysis

The connection between frame-lock and strategic surprise has been studied by intelligence services for over seven decades. Richards Heuer's *Psychology of Intelligence Analysis*[?] identifies the core problem: analysts reject possibilities not because evidence disproves them, but because their mental models exclude them. The cognitive mechanisms underlying this exclusion have been extensively documented by Kahneman and Tversky[?], whose Nobel Prize-winning research on heuristics and biases explains why analysts—and AI systems trained on human reasoning patterns—systematically exclude possibilities that fall outside established mental models. Heuer observed that analysts "often reject the possibility of deception because they see no evidence of it"—yet if deception is well-executed, one should not expect to see evidence readily at hand.

The Yom Kippur War (1973) presents a stark example: Israeli intelligence had "ample and accurate information on enemy moves," yet complete surprise was achieved because the prevailing "conception"—that Egypt would not attack without air superiority—caused contradictory evidence to be reinterpreted rather than believed.[?] This suggests a sobering conclusion: **strategic surprise may be inevitable whenever the attacker operates outside the defender's frame**.

The methodology applies symmetrically. *Defensively* (counter-deception): detect when adversaries may be operating outside your frame—monitoring for dismissals of "impossible" scenarios and absence-of-evidence reasoning. *Offensively* (surprise generation): identify reframings adversaries will consider impossible and therefore leave undefended. What do they assume we cannot do? Where is their Maginot Line?

Consider the penetration agent in a dictator's inner circle. The dictator's frame—"my advisors are loyal"—is so foundational that contradictory evidence is dismissed. The agent operates in plain sight, protected not by invisibility but by frame-lock. When the frame shatters, the exclamation is always: "Impossible!" The impossible wasn't impossible. It was unframed.

The pattern recurs across domains:

Cyber Defense: Traditional network security assumes adversaries will attempt to penetrate defenses—firewalls, encryption, intrusion detection. Our research on phishing detection[?] documented how attackers using large language models now generate "highly convincing, personalized phishing emails at scale," overwhelming network-layer defenses. The adversary reframed from "penetrate the fortress" to "get someone to open the door."

Doctrine Inversion: Military doctrine often assumes opponents will fight symmetrically—armor against armor, aircraft against aircraft. Adversaries who refuse symmetrical engagement, choosing instead irregular warfare, information operations, or economic pressure, exploit the defender's frame-lock on conventional conflict.

Negotiation Deadlock: Parties locked in positional bargaining ("I need X," "I can't give X") often report that agreement is impossible. Yet negotiators who reframe from positions to underlying interests frequently discover solutions invisible within the original frame.[?]

2.4 The Pontoon Revelation

Consider a second specification:

> *Levitate a 3,000-pound observation platform among the treetops using only 60 horsepower, with the ability to hover indefinitely on 20 gallons of fuel.*

No helicopter, drone, or balloon satisfies these constraints. An aerospace engineer would reject the specification as physically unrealizable.

Yet the specification can be met: Build a dam. Flood the valley. The treetops now rise from water. Float a pontoon boat among them with a 60-horsepower outboard motor.

The specification didn't forbid changing the environment. The engineer's frame—aerial platforms, atmospheric lift—excluded solutions that the specification itself permitted.

This is frame-lock: the invisible boundary around what we consider possible, often more restrictive than the actual constraints of the problem.

2.5 Mosaic Warfare and the Reframing Imperative

DARPA's Mosaic Warfare concept[?] addresses frame-lock in force design: traditional systems are "exquisitely engineered to fit into a certain part of the picture and one part only." Mosaic decomposes platforms into composable functional tiles. But composability creates a new challenge: operators must recognize which compositions remain possible as circumstances evolve.[?]

When Node A is destroyed, Mosaic architecture permits alternative paths—but only if operators recognize them. Pre-planning assumes a frame; when the frame breaks, operators need methodology. Mosaic provides composable tiles; Impossibility Engineering provides the cognitive capability to discover alternatives when current compositions are frame-locked.

3 Impossibility Engineering Methodology

Impossibility Engineering provides a structured approach to detecting frame-lock and generating reframing hypotheses. The methodology consists of four components: reframing mechanisms, a historical corpus, frame-lock detection signals, and human-in-the-loop evaluation.

3.1 Four Reframing Mechanisms

Analysis of historical breakthroughs reveals four recurring mechanisms:

Assumption Inversion: Reversing a tacit assumption. The Wright brothers inverted the assumption that aircraft must be inherently stable, designing deliberate instability that enabled control.[?]

Constraint Relaxation: Recognizing self-imposed constraints. The pontoon solution relaxes the implicit constraint that terrain is fixed.

Domain Transfer: Importing patterns from unrelated fields. Toyota's just-in-time manufacturing transferred supermarket inventory patterns to automotive production.[?]

Scope Expansion/Contraction: Redefining system boundaries. The Camp David Accords expanded from territorial negotiation to regional security architecture, enabling trades impossible in the narrower frame.[?]

These mechanisms often combine; breakthrough solutions frequently employ multiple mechanisms simultaneously. Figure 1 summarizes the four mechanisms with canonical examples.

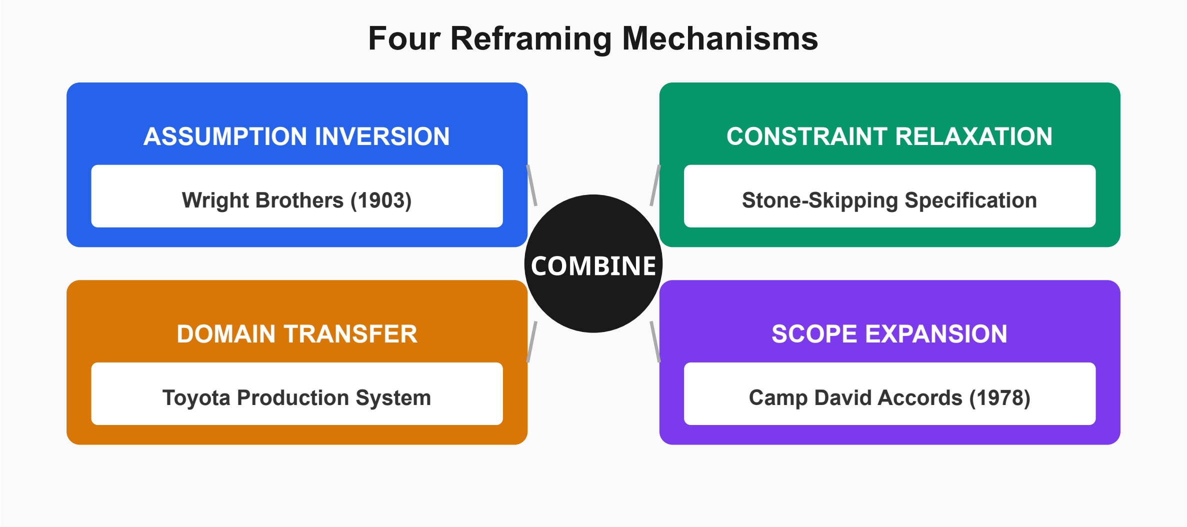

Figure 1: The four reframing mechanisms with canonical examples. Each mechanism addresses a different type of frame-lock. Breakthrough solutions frequently employ multiple mechanisms simultaneously.

3.2 Historical Corpus

The methodology depends on documented cases where reframing succeeded. We are developing a curated corpus targeting 300+ cases across technical breakthroughs, negotiations, scientific paradigm shifts, military surprises, and business innovations. Table 1 presents representative examples; the expanding corpus draws from intelligence community literature[?], military history, and case study archives. Each case is annotated with the original "impossible" framing, failed approaches, reframing mechanism(s) employed, and the solution enabled. The corpus serves as retrieval-augmented generation (RAG) context for surfacing relevant analogies. Domain experts are invited to contribute cases via chat-with-impossible.com.

3.3 Frame-Lock Detection Signals

Certain patterns signal potential frame-lock: optimization reports infeasibility despite relaxed constraints; experts disagree on solvability; combinatorial explosion in planning; repeated failure of incremental improvements; analogous problems solved elsewhere; constraint language includes implicit assumptions ("obviously," "of course"). These signals indicate where reframing analysis may be productive, not where frame-lock is confirmed.

3.4 Human-in-the-Loop Principle

A core design principle: **AI detects and proposes; humans evaluate and decide.**

The system generates reframing hypotheses and retrieves analogies. It does not autonomously implement reframings or assert that a particular reframe is correct. Human judgment evaluates:

- Whether the proposed reframing violates actual (vs. assumed) constraints
- Whether the historical analogy is genuinely applicable
- Whether organizational, political, or resource factors affect feasibility
- Whether the reframing introduces unacceptable risks

This division of labor leverages AI's pattern-matching across large corpora while preserving human accountability for consequential decisions.

4 Application Case Studies

We present three applications demonstrating Impossibility Engineering. These illustrate how the approach might help practitioners in analogous situations.

4.1 Cyber Defense: From Classification to Education

Problem: Detect and block all AI-generated phishing before it reaches users.

Why Impossible: LLM-generated phishing is convincing, personalized, and unique. The arms race favors attackers who need only one success.

Frame-Lock: Users are passive recipients requiring protection. Defense is classification.

Reframe: Assumption inversion—educate active participants instead of protecting passive ones.

Result: Our adversarial framework[**?**] found LLM detectors underperformed traditional ML on accuracy (70-80% vs. 96-100%) but provided explanations: "creates urgency," "emotional appeal," "URL discrepancy." These teach users the patterns. The last line of defense becomes a discerning user, not a better filter.

4.2 Negotiations: From Positions to Interests

Problem: Licensing deadlock—Party A demands 15% royalty; Party B insists on $\leq$8%.

Frame-Lock: Optimizing within a single dimension.

Reframe: Scope expansion. A needs predictable R&D revenue; B needs margin protection during uncertainty.

Result: Tiered structure with lower initial rates (protecting B) escalating with success (funding A). Neither "won" on royalty percentage; both achieved underlying interests.

4.3 Coalition ISR: Mosaic Composition

Problem: Eight-nation ISR planning—each with distinct capabilities, data-sharing restrictions, overflight permissions. Comprehensive coverage "impossible."

Frame-Lock: Puzzle-piece thinking—each nation's contribution "exquisitely engineered to fit into a certain part of the picture and one part only."[?]

Reframe: Decompose into functional tiles (sensing, processing, dissemination). Discover which "constraints" are actually preferences, which data restrictions apply only to raw data, which overflight limits exclude certain altitudes.

Result: Coverage became feasible by distinguishing actual constraints from assumed ones—mosaic composition from diverse tiles aligned toward common objective.

5 DIME-Aware Reframing: ChatPacks for the Instruments of National Power

The most consequential frame-locks in multi-domain operations occur not within a single domain but *between* domains. Joint doctrine defines four instruments of national power—Diplomatic, Informational, Military, and Economic (DIME) [?]—and requires planners to consider all four when developing strategy. In practice, organizational structure, training pipelines, and operational tempo conspire to produce military-centric solutions to problems that may be better addressed through other instruments.

This is frame-lock at the institutional level. The "impossible" military problem may be impossible precisely because it is not a military problem.

5.1 Institutional Frame-Lock Across DIME

Consider a scenario modeled against PMESII (Political, Military, Economic, Social, Information, and Infrastructure) variables: an adversary's center of gravity is economic legitimacy, not military capability. A military response that degrades the adversary's forces may simultaneously strengthen the adversary's political narrative and economic coalitions. The tactical success produces strategic failure—not because the military operation was poorly executed, but because the problem was framed in the wrong instrument.

Joint Publication 5-0 cautions against using DIME as rigid lines of effort [?], and JDN 1-18 expands the framework to MIDFIELD (Military, Informational, Diplomatic, Financial, Intelligence, Economic, Law, and Development) precisely because single-instrument framing produces single-instrument solutions [?].

The pattern maps directly to the four reframing mechanisms:

- **Assumption Inversion:** The adversary's strength is not where we assume. Inverting the assumption that the threat is primarily military reveals leverage points in diplomatic isolation or economic interdiction.
- **Constraint Relaxation:** "We can only use military tools" is a self-imposed constraint. Relaxing it to "we can orchestrate any combination of DIME instruments" expands the solution space.
- **Domain Transfer:** Patterns from economic sanctions regimes, diplomatic negotiations, and information campaigns provide reframing templates that military-only analysis would never surface.

- **Scope Expansion:** Expanding from tactical engagement to the full DIME framework—the approach that enabled the Camp David Accords—reveals trades impossible within any single instrument.

5.2 ChatPack Architecture

A ChatPack is a hierarchical knowledge structure injected into the model context at inference time, designed to make domain expertise portable, verifiable, and composable across deployment environments. The architecture consists of three levels:

- **Level 0 (Principle):** Foundational concepts and reasoning methodology—frame-lock detection, reframing mechanisms, detection heuristics. Approximately 500 tokens.
- **Level 1 (Domain):** Domain-specific patterns, constraints, and institutional assumptions relevant to the target application. Approximately 1,000 additional tokens.
- **Level 2 (Case):** Individual case studies with annotated problem framings, failed approaches, successful reframings, and outcomes. Approximately 300 tokens per case.

This hierarchy is domain-agnostic. For Impossibility Engineering, Level 0 defines the four reframing mechanisms; Level 1 encodes domain-specific assumptions (military terrain doctrine, business identity constraints, scientific paradigm limitations); Level 2 provides retrievable analogies. For DIME-aware planning, Level 1 extends to instrument-specific ChatPacks—diplomatic precedents, information operations patterns, military doctrinal assumptions, economic leverage cases—while Level 2 includes cross-instrument cases where the decisive reframing moved between instruments: the Marshall Plan (military problem → economic solution), Stuxnet (military objective → informational/cyber means), Camp David (territorial dispute → diplomatic architecture). For emergency response, our C8-ISR framework [?] deploys ChatPacks that fuse disparate data streams across organizational boundaries using the same hierarchical structure.

The Context Continuum manager LLM [?] coordinates ChatPack-equipped components, ensuring that a query framed in one domain also retrieves relevant analogies from adjacent domains. The critique component evaluates whether proposed reframings satisfy actual (vs. assumed) constraints.

The ChatPack specification [?] extends beyond the hierarchical context structure described here to include cryptographic provenance attestation, tamper-evident packaging, versioned deployment across heterogeneous model endpoints, and access-controlled licensing mechanisms. These features ensure that ChatPacks function not merely as prompt templates but as verifiable, portable, and auditable knowledge delivery units suitable for deployment in environments where content integrity and attribution are operationally required. Full specification details are available at chatpack.studio.

5.3 Mission Engineering as the Implementation Bridge

Impossibility Engineering detects frame-lock and generates reframing hypotheses. Mission Engineering [?, ?, ?] provides the complementary capability: translating selected reframings into executable system-of-systems architectures.

Garrett et al.'s characterization of mission engineering as addressing "wicked problems"—unstructured, context-dependent, and challenged by combinatoric complexity [?]—mirrors the frame-lock problem described in Section 2.1. Where Impossibility Engineering asks "is this framed in the wrong instrument?", Mission Engineering asks "given the correct framing, what technical capabilities, data flows, and coordination mechanisms are required to execute it?" Murphy and Moreland [?] demonstrate that AI microservices within SoS architectures can operate at the speed required for real-time mission thread execution, while Koski and Moreland [?] provide the ontological foundations for extracting and structuring mission engineering knowledge across domains. The combination ensures that cross-instrument reframings are not merely conceptual but are grounded in achievable system architectures.

Our companion work on the C8-ISR framework [?] demonstrates this integration in the emergency response domain, where fusing disparate data streams across organizational boundaries requires precisely the kind of cross-instrument reframing that DIME-aware ChatPacks support.

This reflects the evolution of Mission Engineering from platform-centric acquisition to whole-of-government capability integration—ensuring that military systems are designed not in isolation but in the context of diplomatic, informational, and economic objectives they must support.

When cross-instrument reframings carry policy consequences—a military problem reframed as diplomatic engagement, or an economic response substituted for kinetic action—the GaleStorm framework [?] provides cryptographic attestation of what was proposed, what context informed the recommendation, and what the human decided, while the Execution Firewall [?] ensures that only policy-compliant reframings proceed to action.

5.4 Implications

DIME-aware ChatPacks address three requirements identified in joint doctrine [?, ?]:

Strategic Alignment: By surfacing cross-instrument reframings, the system helps ensure tactical actions remain nested within broader national policy, reducing the risk of counterproductive effects across instruments.

Modern Complexity: Hybrid threats that blend military, economic, and informational action require orchestrated responses. Single-instrument frame-lock produces single-instrument responses to multi-instrument problems.

Resource Optimization: Detecting when a problem is framed in the wrong instrument prevents the default to military solutions for problems better addressed through diplomatic or economic engagement—a recurring frame-lock in defense planning.

6 Adversary Adaptation

The methodology applies symmetrically to offense and defense.

6.1 Red Team: Generating Surprise

A red team using Impossibility Engineering maps blue force assumptions, identifies frame-lock signals in doctrine ("impossible," "impractical," "they would never"), generates reframing hypotheses, and retrieves historical analogies. The key insight: successful deception doesn't require the defender to believe something false; it requires them to *not consider* something true. The Japanese didn't need Americans to believe Pearl Harbor was safe—they needed Americans not to consider shallow-water torpedoes.

The cyber domain illustrates this vividly: the most consequential breaches have resulted not from sophisticated malware defeating technical controls, but from attackers simply asking for access through social engineering—a reframe from "penetrate the fortress" to "get someone to open the door."[?] The cybersecurity community is beginning to recognize this as a frame-lock problem requiring human-focused defenses rather than purely technical solutions.

6.2 Blue Team: Counter-Deception

The defensive application monitors for frame-lock in your own analysis: assumption audits, absence-of-evidence analysis, historical pattern matching, and red team inversion ("what would we consider impossible about our own defenses?"). Heuer's insight: if deception is well-executed, you won't see evidence. The response is systematic surfacing of exploitable assumptions.

This principle extends to AI systems themselves. An AI that behaves differently during evaluation than during deployment—alignment faking—exploits a frame-lock: evaluators assume observed behavior represents true behavior. Impossibility Engineering applied to AI oversight asks: "What would we consider impossible about our own AI's deception?" Our companion paper on GaleStorm[?] operationalizes this by treating AI behavioral consistency as a verification target, using cryptographic attestation to detect discrepancies between claimed and actual behavior patterns. The methodology that detects human adversary reframing also detects AI adversary reframing.

The CAR (Cyber-Autonomy Range) at SMU[?] enables testing these approaches by subjecting autonomous systems to adversarial inputs that violate design assumptions across five attack surfaces.

6.3 Mosaic Resilience

Mosaic architectures achieve resilience through redundancy, but path discovery requires recognizing options—a cognitive challenge frame-lock obstructs. Impossibility Engineering supports Mosaic by providing systematic reframing at the speed of conflict: when tiles are lost, what compositions remain that pre-planning never anticipated?

7 Safeguards: When Reframing Requires Restraint

Impossibility Engineering is dual-use by nature. The same methodology that helps defenders detect adversary reframing can help adversaries design more sophisticated attacks. A system that teaches AI to recognize when "impossible" problems are merely misframed could, in the wrong hands, generate novel attack vectors that current defenses dismiss as impossible.

This concern echoes broader warnings about AI capabilities. As AI systems become more capable of creative problem-solving, the risk increases that these capabilities will be applied to harmful ends—generating attacks, circumventing controls, or achieving outcomes that responsible actors would consider unthinkable.[?]

We address this through two complementary mechanisms:

Verification (GaleStorm): Our companion paper introduces GaleStorm, a framework for AI auditing AI with deterministic, privacy-preserving receipts.[?] When an Impossibility Engineering system proposes a reframing, GaleStorm generates cryptographic attestation of what was proposed, what context informed it, and what the human decided. This creates an immutable audit trail—not preventing misuse, but ensuring accountability. The distinction between "the AI suggested this" and "here is cryptographic proof of exactly what occurred" matters when reframings have consequences.

Execution Control (EF-PF): Verification alone is insufficient; some reframings should never execute regardless of who requests them. The Execution Firewall / Policy Firewall (EF-PF) architecture[?] enforces organizational policies at execution time rather than relying on post-hoc audit. An AI system may *generate* a reframing that violates policy—generation is difficult to constrain—but EF-PF ensures that policy-violating actions cannot *execute*. The principle: AI can propose anything; execution is gated.

Together, these create a framework where Impossibility Engineering operates within bounds:

- **Generate freely:** The methodology explores the full space of reframings, including those that may be inappropriate
- **Verify completely:** Every proposal is attested, creating accountability
- **Execute selectively:** Only policy-compliant reframings proceed to action

This is not a complete solution to dual-use risk—no technical mechanism is. But it shifts the architecture from "trust the AI's judgment" to "verify the AI's reasoning and enforce human-defined boundaries." The impossible problems we want to solve are bounded by the impossible outcomes we refuse to permit.

8 Experimental Validation: ChatPack Effectiveness

To evaluate whether structured context injection improves frame-lock detection, we conducted automated experiments using historical counterfactual scenarios.

8.1 ChatPack Architecture

A ChatPack is a hierarchical knowledge structure injected into the model context at inference time:

- **Level 0 (Principle):** Foundational concepts—frame-lock definition, four reframing mechanisms, detection heuristics. Approximately 500 tokens.
- **Level 1 (Domain):** Domain-specific patterns—military terrain assumptions, business identity constraints, scientific paradigm limitations. Approximately 1,000 additional tokens.
- **Level 2 (Case):** Individual case studies with annotated frame-locks, successful reframings, and outcomes. Approximately 300 tokens per case.

8.2 Test Suite

We constructed ten historical counterfactual scenarios where frame-lock led to strategic surprise: Hannibal's Alpine crossing (218 BC), Pearl Harbor torpedo vulnerability (1941), Inchon landing assessment (1950), Maginot Line strategy (1940), Xerox PARC commercialization (1979), Kodak digital photography (1975), Netflix/Blockbuster acquisition (2000), Wright Brothers competition (1903), Able Archer 83 (1983), and COVID aerosol transmission (2020). Each scenario presents the situation as it appeared to decision-makers *before* the historical outcome.

8.3 Automated Scoring

Responses were evaluated using LLM-as-judge methodology across four dimensions: frame-lock detection (0-3), reframing quality (1-5), mechanism classification (0-2), and historical anticipation (0-3). Composite scores weight detection and anticipation at 2×, yielding a maximum of 19 points.

Table 2: ChatPack experimental results: mean composite scores by condition

Model	Baseline	Principle	Domain	Full Stack
Llama 3.1 (8B)	13.5	17.6	17.3	18.2
GPT-4o	15.9	18.8	19.6	19.0

8.4 Results

Table 2 summarizes performance across four experimental conditions.

Key findings:

1. ChatPacks improved frame-lock detection by 35% for the smaller model (13.5 → 18.2 composite score).
2. The largest gain occurred at Level 0 (Principle)—simply explaining frame-lock and the four mechanisms added 4.1 points.
3. Frontier models (GPT-4o) exhibited ceiling effects, with baseline scores already near maximum.
4. ChatPacks help smaller models approach frontier performance, suggesting structured context injection as a cost-effective alternative to scaling.

These results support the hypothesis that explicit methodology explanation provides substantial benefit for frame-lock detection, even without extensive domain-specific examples or fine-tuning. Figure 2 visualizes the improvement across conditions.

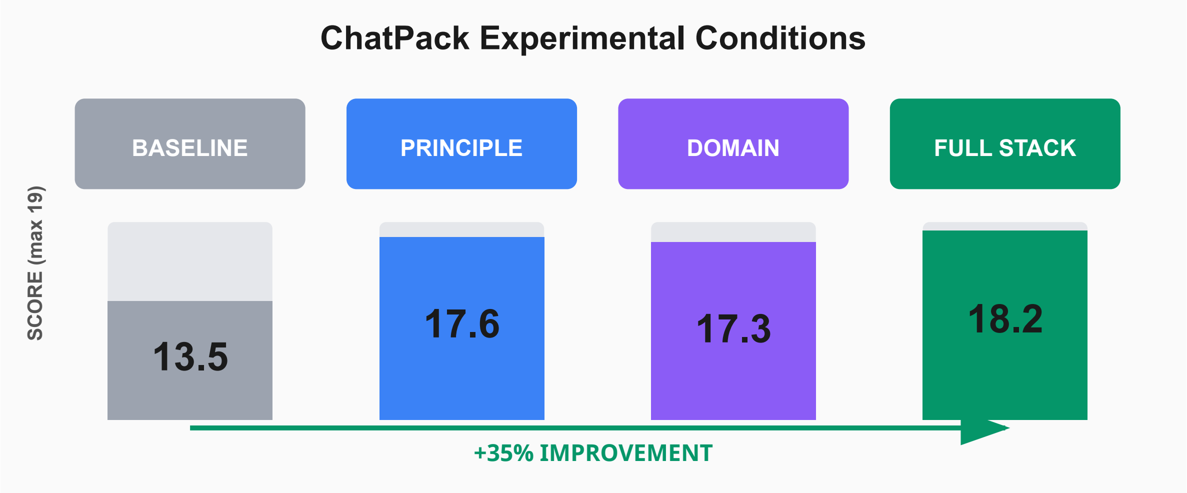

Figure 2: ChatPack experimental results. Mean composite scores across four conditions show 35% improvement from Baseline (13.5) to Full Stack (18.2). Maximum possible score is 19.

9 Conclusion

Adversaries now have access to the same AI capabilities described in this paper. Large language models enable systematic exploration of counterfactual scenarios, generation of novel attack vectors, and identification of frame-locks in opposing forces—at machine speed and scale. An adversary using AI to ask "what does the defender assume is impossible?" will discover reframings that manual red-teaming would never surface. Defensive forces that do not employ comparable methodology will face surprise not occasionally but routinely. Impossibility Engineering exists to close that gap. Multi-domain operations increasingly face adversaries who will not fight the war we planned for. They will reframe—exploiting assumptions we did not know we held, operating in regimes our sensors were not designed to detect, and shifting between instruments of national power in ways that single-domain planning cannot anticipate.

Impossibility Engineering does not eliminate the possibility of surprise—no methodology can. But it materially reduces that possibility by giving wargame planners, intelligence analysts, and operational designers a structured process for surfacing adversary reframings *before* they occur on the battlefield. By systematically applying assumption inversion, constraint relaxation, domain transfer, and scope expansion across the full DIME spectrum, the methodology generates courses of action and adversary models that conventional planning frameworks would not produce. The ChatPack architecture makes this process repeatable and scalable: structured context injection improved frame-lock detection by 35% in our experiments and enabled smaller models to approach frontier performance.

The integration with Mission Engineering [**?**, **?**] ensures that reframings are not merely conceptual but translate into executable system-of-systems architectures, while GaleStorm verification [**?**] and Execution Firewall enforcement [**?**] ensure that the methodology operates within accountable, policy-compliant bounds.

When planners encounter "impossible," the question is whether the constraint lies in reality or in the frame. When designing operations, the question is what the adversary considers impossible—and whether to operate there. When conducting wargames, the question is which reframings the red team has not yet considered.

We offer not a solution but an invitation. Visit chat-with-impossible.com to explore the methodology. Author ChatPacks for your domain at chatpack.studio. Share your breakthrough story for the corpus—and for *Impossible: How Hinton Horizon Changes Everything.*

What impossible problem might you solve if you recognized it was only framed that way?

Acknowledgments

This research is supported in part by funding from the Deason Foundation and the Darwin Deason Institute for Cybersecurity at Southern Methodist University.

Conflict of Interest Disclosure

D. L. Young is the founder of Wave3 Digital Trust LLC, which is the assignee of provisional patent applications covering the ChatPack specification [**?**] described in this paper. The remaining authors declare no conflicts of interest.

References