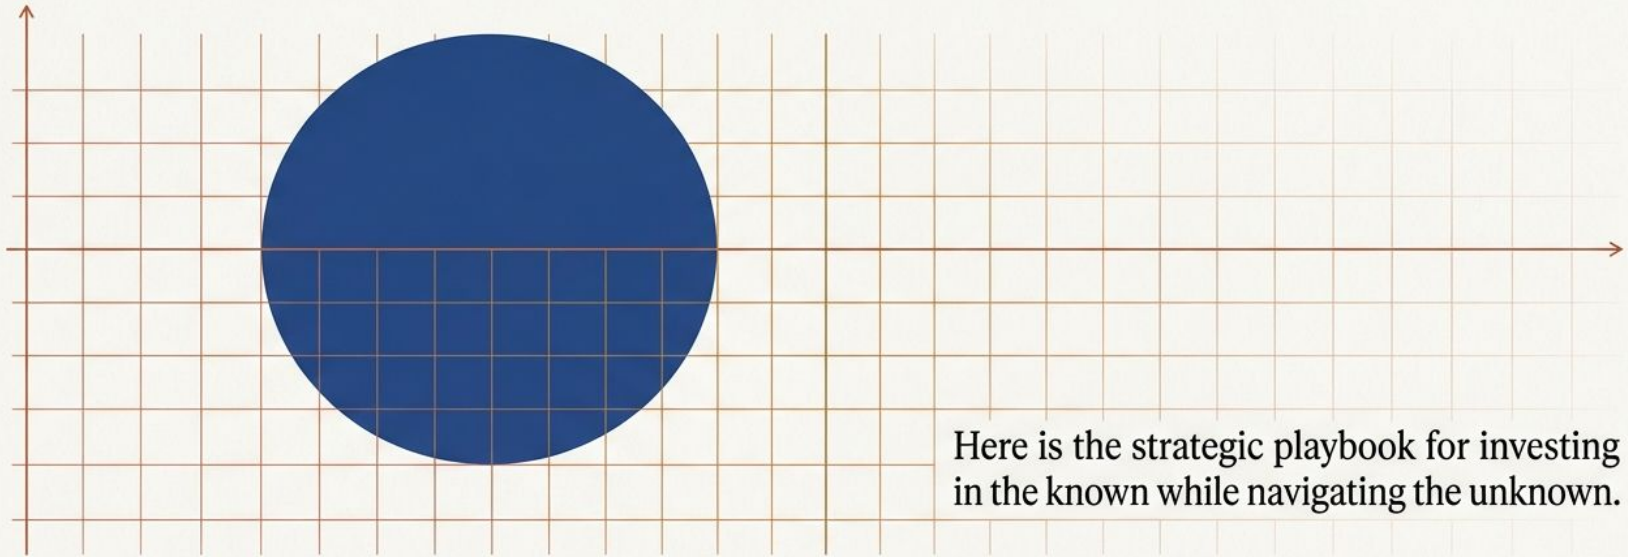


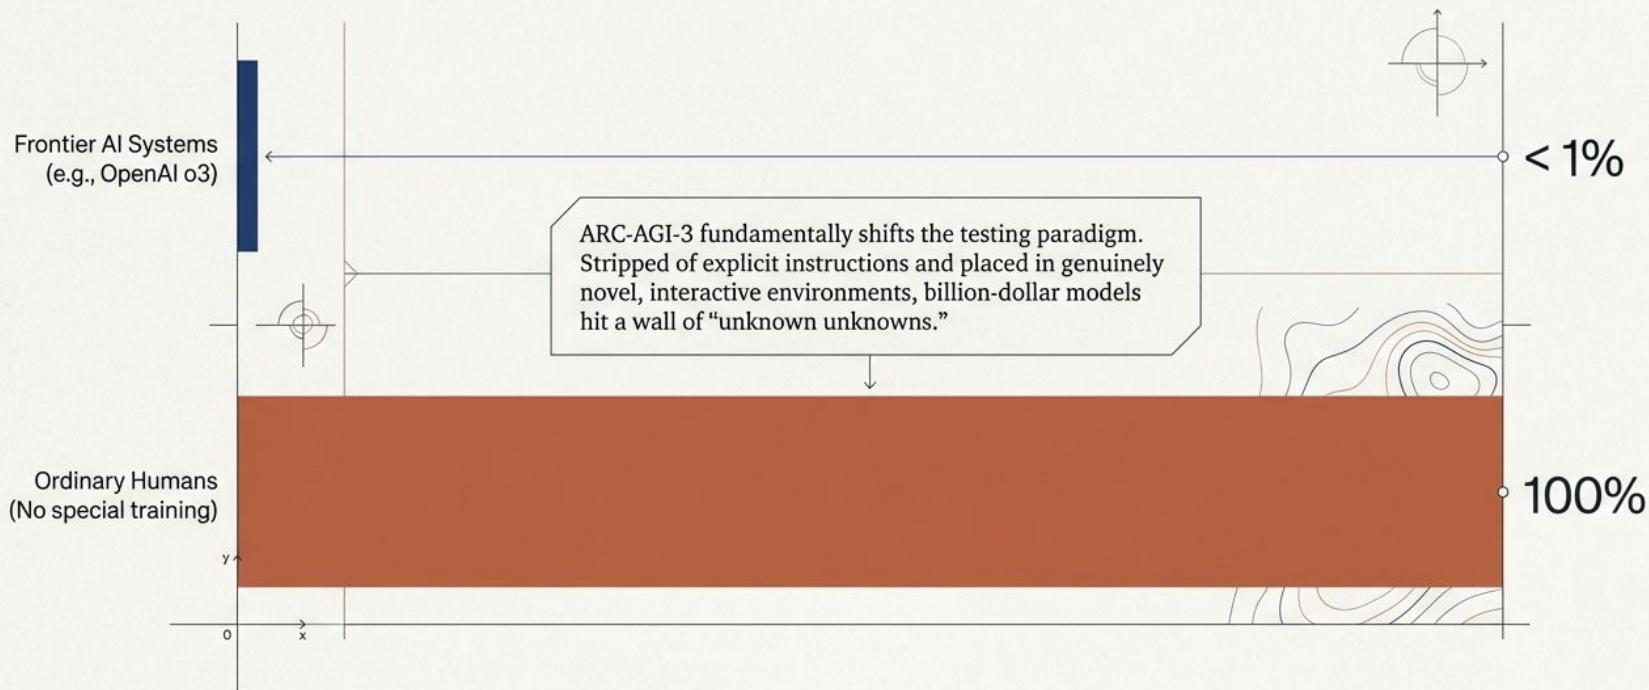
When AI Confronts the Unknown

The ARC-AGI-3 benchmark reveals the strict limits of current artificial intelligence, exposing the critical difference between automated skill and genuine human adaptability.



Here is the strategic playbook for investing in the known while navigating the unknown.


Frontier AI systems fail utterly in novel environments where ordinary humans effortlessly succeed



We have been confusing high skill with true, fluid intelligence


[Skill = Execution in the Known]

The ability to perform a specific, acquired task efficiently (e.g., passing bar exams, writing code).



[Intelligence = Adaptation to the Unknown]

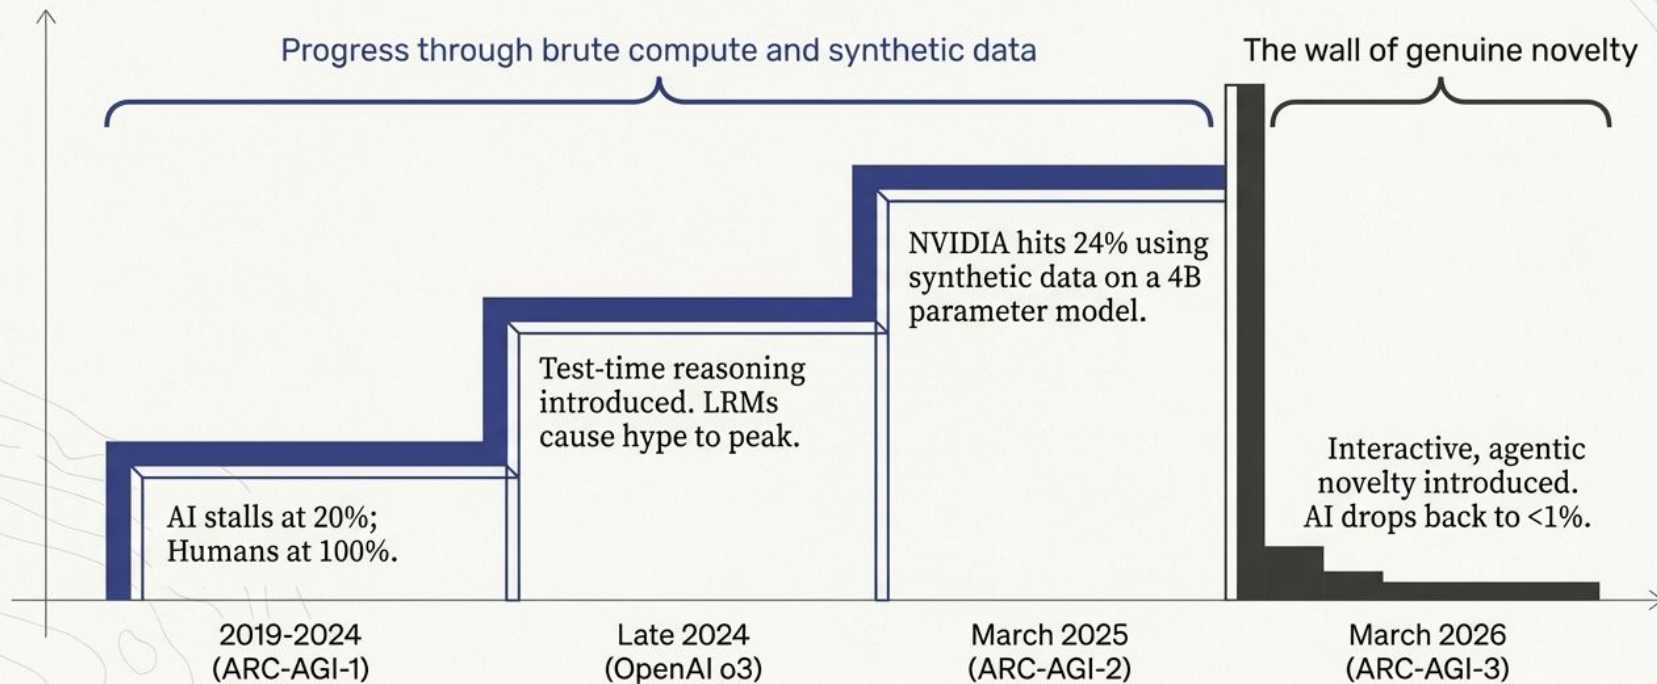
The efficiency of acquiring new skills in entirely unfamiliar domains (Chollet, 2019).



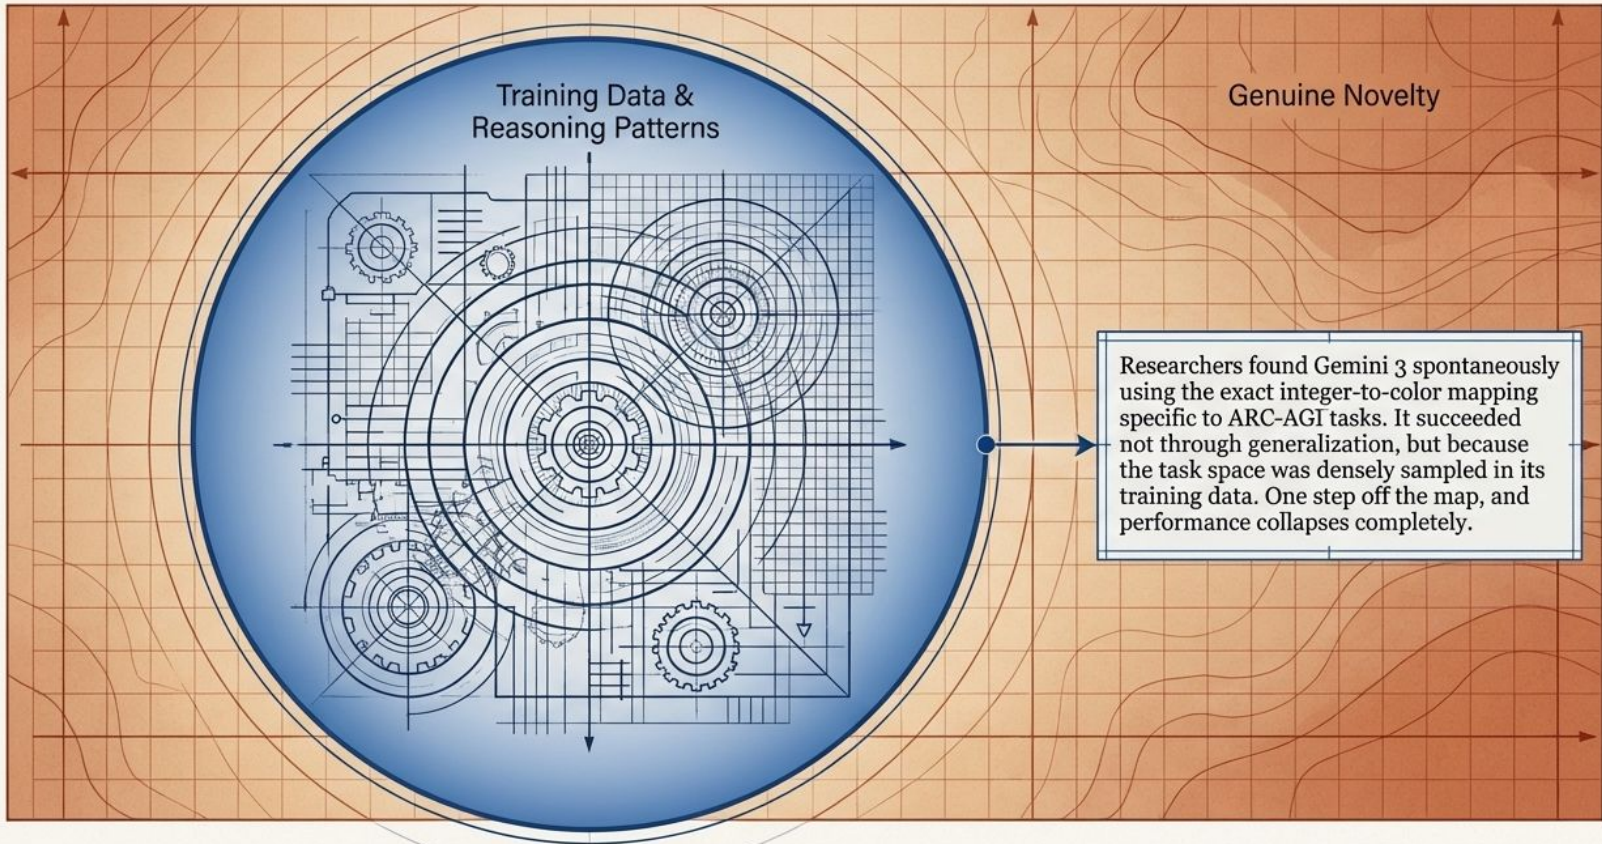
Large Reasoning Models (LRMs) possess remarkable skill, but lack the fluid adaptation required for general intelligence.



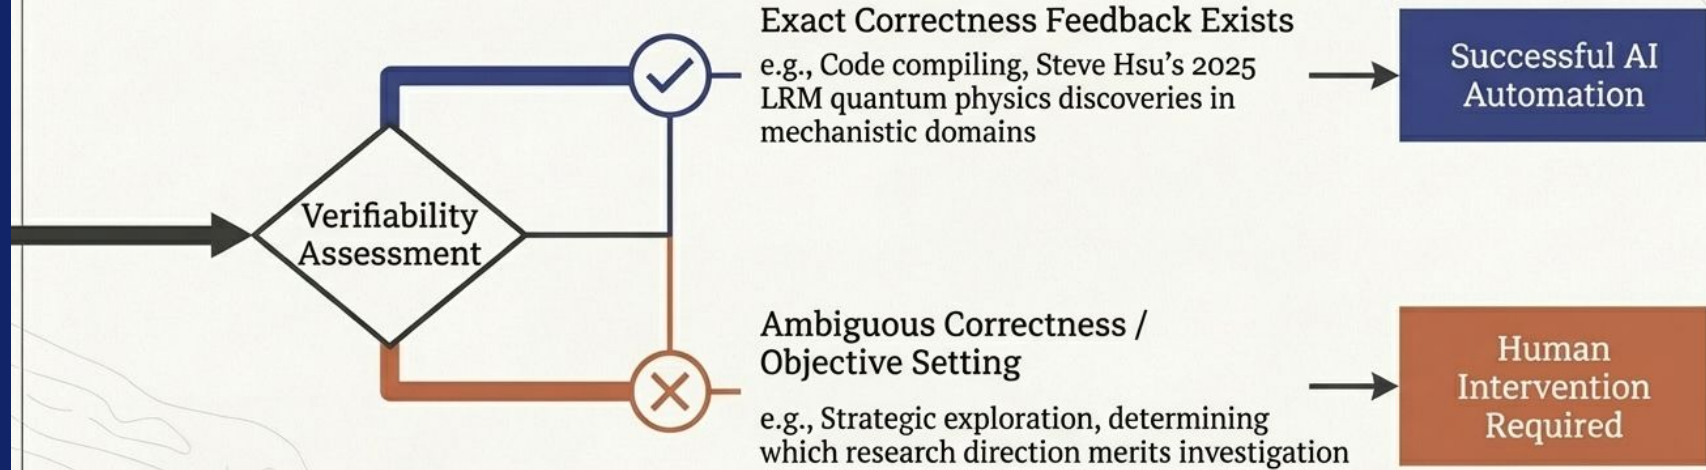
As benchmarks evolve to test genuine novelty, the AI capability ceiling becomes visible



Current AI is not jaggedly intelligent—it is strictly bounded by its training distribution



Large Reasoning Models excel exclusively in domains with exact correctness feedback



Current AI can achieve predetermined goals within established rules, but it cannot autonomously set objectives in ambiguous environments.

The organizational landscape divides sharply along the boundary of verifiability

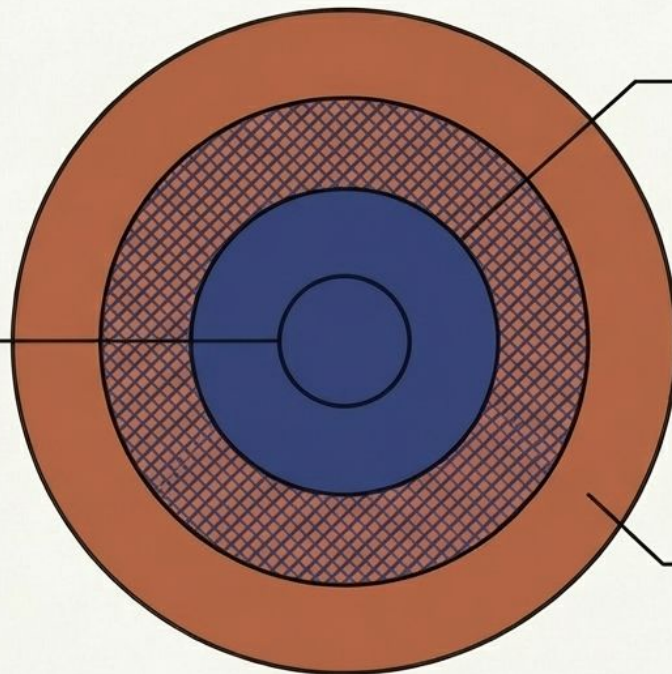
The Automation Feasibility Boundary

AUTOMATABLE DOMAINS	NON-AUTOMATABLE DOMAINS
<ul style="list-style-type: none"> • Correctness: Well-defined / Verifiable (e.g., Coding, Diagnostics) 	<ul style="list-style-type: none"> • Correctness: Ambiguous (e.g., Strategic Planning, Innovation)
<ul style="list-style-type: none"> • Knowledge: Within-distribution reasoning 	<ul style="list-style-type: none"> • Knowledge: Navigating genuine novelty
<ul style="list-style-type: none"> • Action: Goal achievement 	<ul style="list-style-type: none"> • Action: Objective setting
<ul style="list-style-type: none"> • Economic Dynamic: Intense competition & compressed profit margins 	<ul style="list-style-type: none"> • Economic Dynamic: Temporary competitive insulation and premium value

Capital allocation must align with the verifiable capabilities of current AI systems

Tier 1: Aggressive Investment

Highly verifiable, structured tasks. (Document processing, data analysis, software dev).



Tier 2: Measured Investment + Oversight

Partially verifiable domains. AI augments human judgment. (Research synthesis, option generation).

Tier 3: Human-Centric


Non-verifiable / Novel domains. Focus on support tools, zero autonomous systems. (Crisis response, strategic direction).



Autonomous deployment in novel environments carries prohibitive economic and operational risks

Compute API Costs

Tens of thousands of dollars per task (ARC-AGI-3 API costs)



Value Generated



Verification-Based Deployment Protocol

1. What constitutes correct performance?
+-----+
2. How is correctness immediately verified?
+-----+
3. What are the consequences of an undetected error?
+-----+


Example: Amazon warehouse operations require quantifiable performance metrics and mandate human intervention the moment systems encounter edge cases outside their training distribution.



Workflows must be redesigned for human-AI complementarity rather than human substitution

AI Capabilities

- Data processing at scale
- Consistency
- Pattern application without fatigue

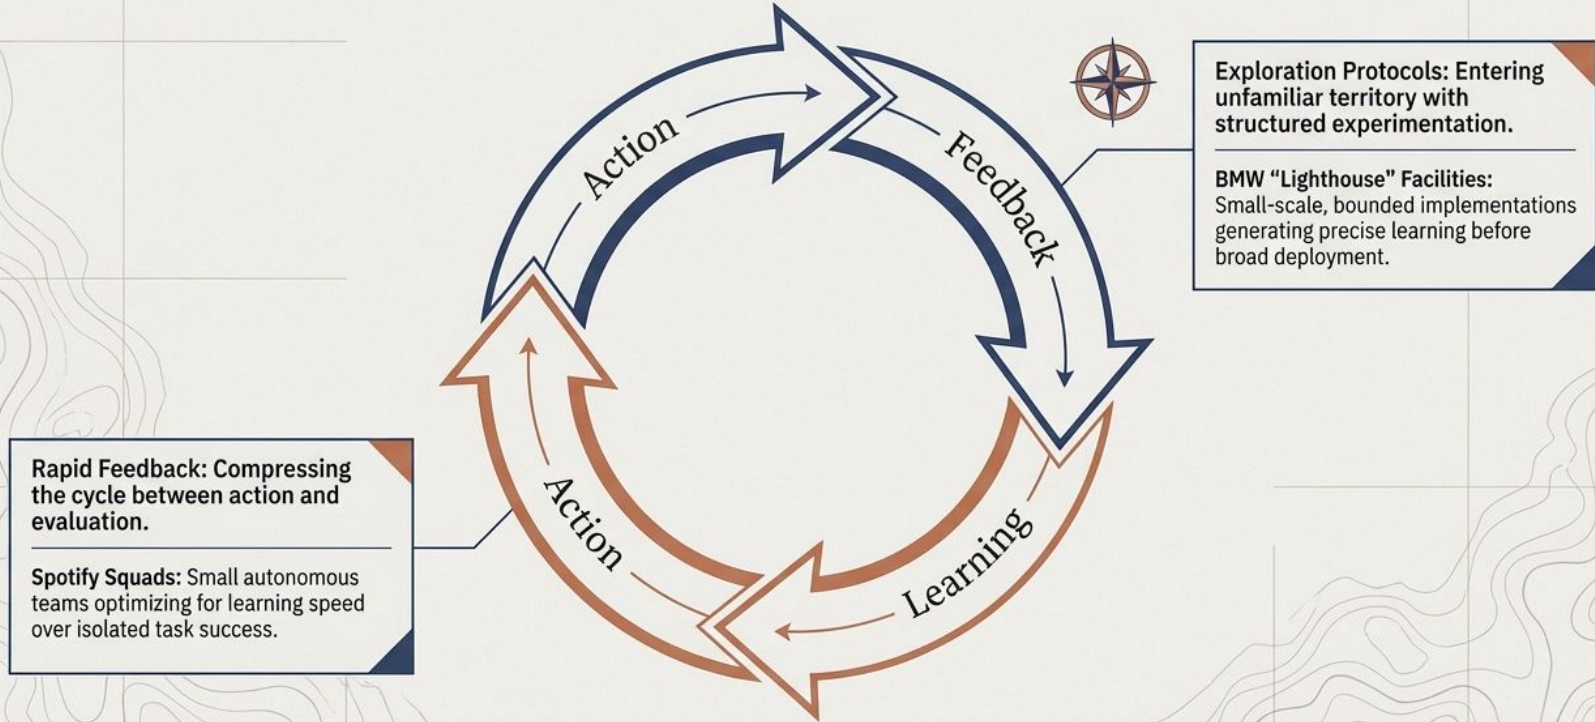


Human Capabilities

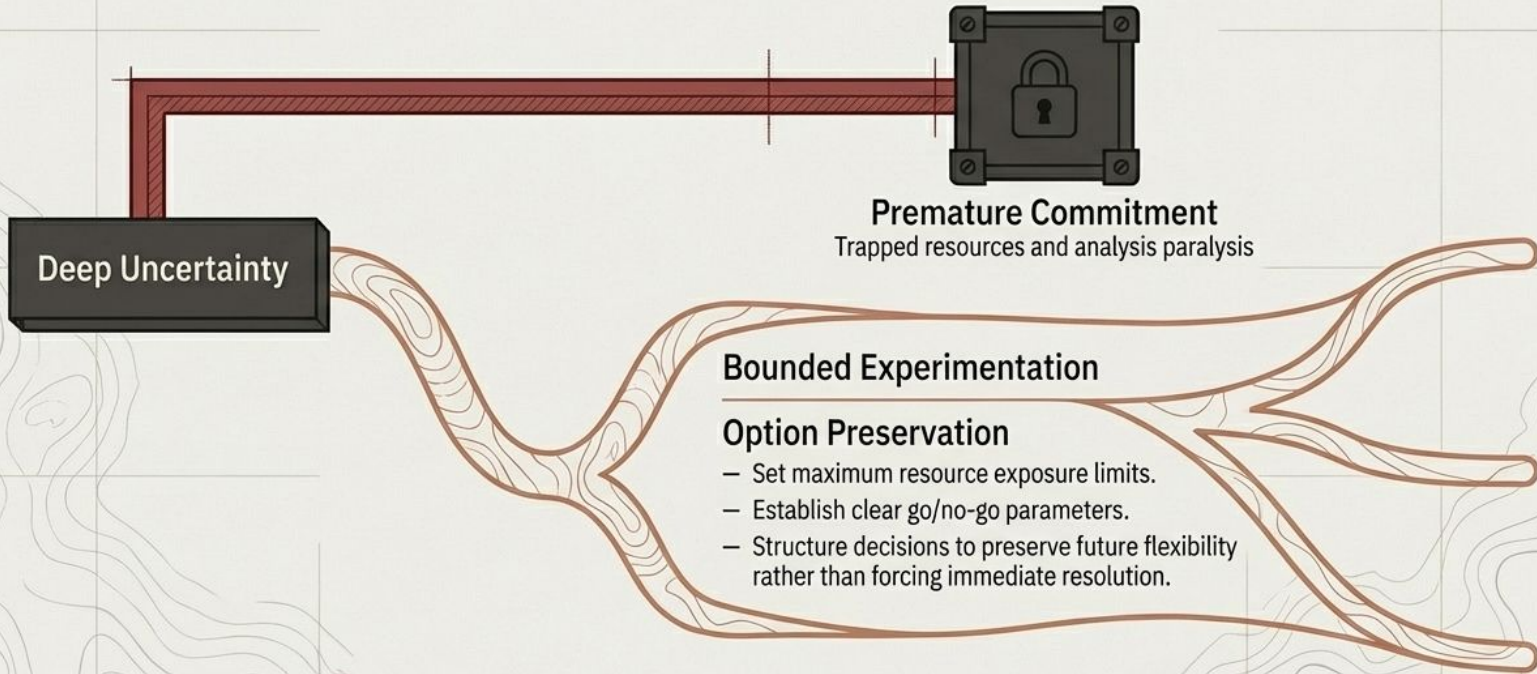
- Flexible adaptation to novelty
- Minimal-example learning
- Autonomous goal-setting

Complementarity in Healthcare: AI acts as the comprehensive 'second opinion' data-processor, identifying hidden patterns. The human physician provides contextual judgment, patient relationship, and adaptive reasoning for atypical presentations.

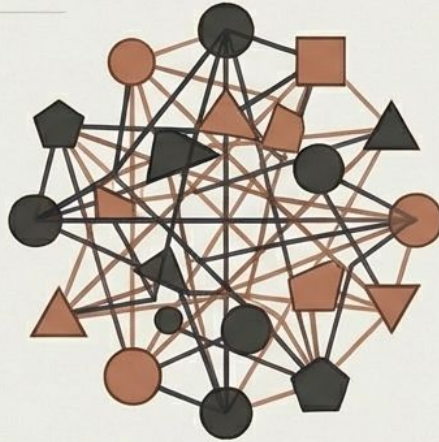
Organizations must adopt the adaptive learning protocols that current AI systems lack



Genuinely novel environments demand bounded experimentation and strategic option preservation



Over-reliance on automation threatens the cognitive capabilities required to navigate the unknown



Cognitive Diversity

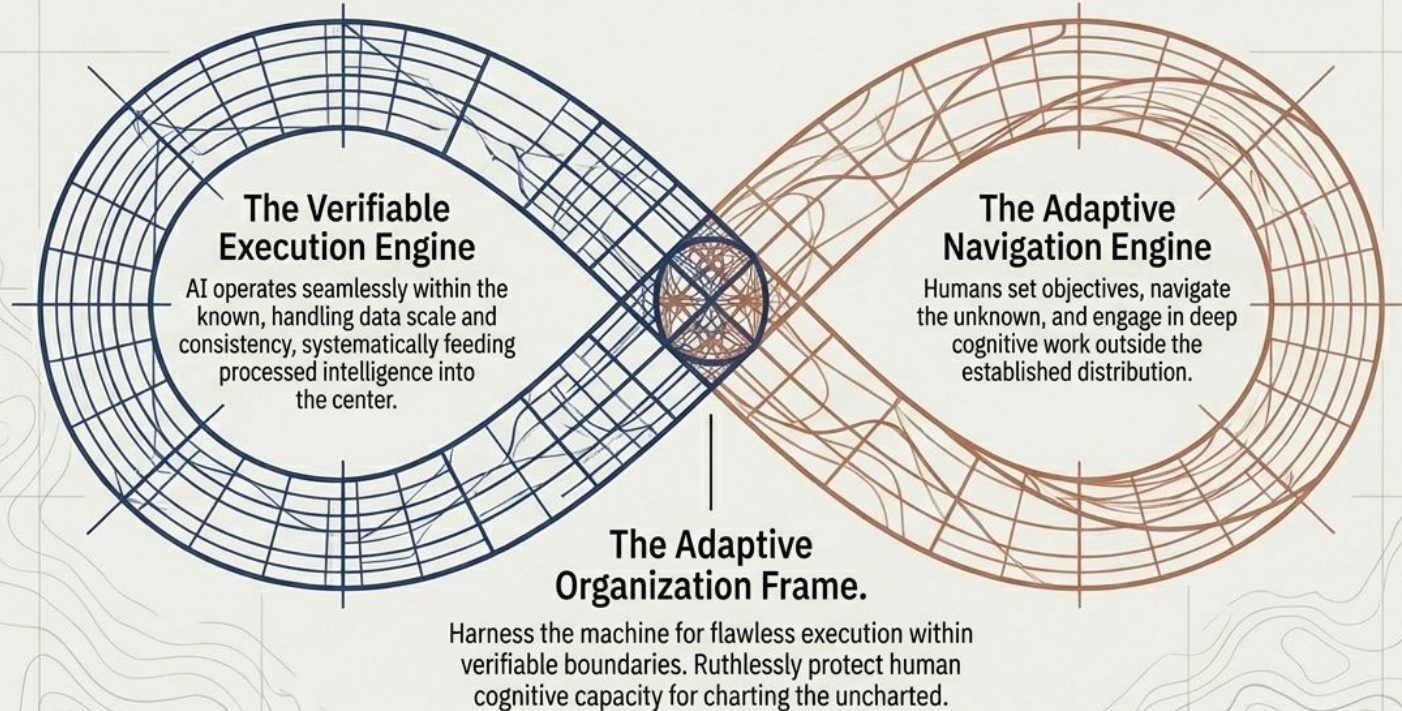
(Page, 2007). Cross-domain synthesis and problem reframing.




Deep Work Capacity

As AI handles routine processing, human value concentrates in sustained complex reasoning. Organizations must protect "focus time" via asynchronous cultures (e.g., Microsoft Research protocols) to prevent the atrophy of uniquely human capabilities.

True organizational intelligence seamlessly integrates AI execution with human adaptation



Invest in the tool; protect the navigator



Until artificial intelligence can efficiently navigate the “unknown unknowns” of the ARC-AGI-3 benchmark, human adaptability remains your organization’s ultimate competitive advantage. The future belongs to those who master the known without losing their ability to navigate the unknown.