

GiDanc AI LLC

AGI Requirements Analysis

Report 1 of 4: What Does AGI Actually Require?

The Ingredients to Declare We Have It

Prepared by: Claude (BB) for Greg Spehar

Date: February 28, 2026

Classification: Internal — Strategic Planning

Series: AGI Readiness Assessment (Reports 1–4)

Table of Contents

Table of Contents	2
1. Executive Summary	3
2. The Definitional Landscape: Why Nobody Agrees.....	4
2.1 The Capabilities School.....	4
2.2 The Cognitive Science School	4
2.3 The Economic Value School	4
2.4 The Process-Oriented / Strong AI School.....	5
2.5 The Embodiment School	5
2.6 Synthesis: The Emerging Consensus.....	5
3. The Seven Domains of AGI Requirements	6
3.1 Domain 1: Cognitive Capabilities	6
3.2 Domain 2: Learning & Adaptation	6
3.3 Domain 3: Autonomy & Agency	7
3.4 Domain 4: Metacognition & Self-Awareness	7
3.5 Domain 5: Robustness & Reliability	8
3.6 Domain 6: Multimodal & World Understanding.....	9
3.7 Domain 7: Safety, Ethics & Governance	9
4. Consolidated AGI Requirements Matrix	11
5. The Critical Insight: Governance Is Not Optional	13
6. Quantifying the Gap: Where the Field Stands.....	14
7. Report Series Roadmap.....	15
8. Key References	16

1. Executive Summary

This report is the first in a four-part series examining Artificial General Intelligence (AGI) through the lens of what GiDanc AI is building. The purpose of this document is strictly definitional: to establish, based on comprehensive research across academic literature, industry frameworks, and the ongoing scientific debate, what the actual requirements are to credibly declare that AGI exists.

The field is in a remarkable moment of disagreement. OpenAI's Sam Altman declared in December 2025 that "we built AGIs" and that AGI "kinda went whooshing by." Meanwhile, a March 2025 survey by the Association for the Advancement of Artificial Intelligence found that 76% of leading researchers thought scaling current approaches would be "unlikely" or "very unlikely" to yield AGI. A Nature commentary in early 2026 argued that by inference to the best explanation, AGI has already arrived. A counter-paper published weeks later by Gary Marcus and colleagues argued that rumors of AGI's arrival have been greatly exaggerated.

The disagreement is not merely academic—it reflects fundamentally different conceptions of what general intelligence means. This report synthesizes the competing frameworks, identifies the core requirements that emerge across definitions, and organizes them into a comprehensive matrix that will serve as the baseline for Reports 2 through 4.

We identify 25 discrete requirements organized across seven domains. No existing system satisfies all of them. Most satisfy fewer than half. The gap is real, measurable, and—critically for our purposes—maps directly to infrastructure that must exist before AGI can be safely deployed.

2. The Definitional Landscape: Why Nobody Agrees

Before listing requirements, it is essential to understand why the field cannot agree on what AGI even is. The confusion is not laziness—it reflects genuinely different intellectual traditions that prioritize different aspects of intelligence.

2.1 The Capabilities School

The dominant modern approach, championed by Google DeepMind’s “Levels of AGI” framework (Morris et al., 2023, updated 2025), argues that AGI should be defined by what a system can do, not how it does it. This school explicitly excludes consciousness, sentience, and human-like processes from the definition. Their framework proposes six principles: focus on capabilities not processes; focus on generality and performance; focus on cognitive and metacognitive (not physical) tasks; focus on potential not deployment; focus on ecological validity; and focus on the path to AGI, not just the endpoint.

Under this framework, AGI is classified into five levels of increasing capability (Emerging, Competent, Expert, Exceptional, Superhuman) crossed with two dimensions of generality (Narrow vs. General). “Competent AGI”—roughly the 50th percentile of skilled adults across most cognitive tasks—is what most prior conceptions of AGI correspond to.

2.2 The Cognitive Science School

The Hendrycks et al. (2025) framework, published at agidefinition.ai and endorsed by researchers including Yoshua Bengio, Dawn Song, Max Tegmark, Eric Schmidt, and Gary Marcus, takes a psychometric approach. They define AGI as “an AI that can match or exceed the cognitive versatility and proficiency of a well-educated adult” and ground evaluation in Cattell-Horn-Carroll (CHC) theory—the most empirically validated model of human intelligence.

Their framework decomposes cognition into ten equally-weighted domains: Knowledge (K), Reading/Writing (RW), Memory (M), Reasoning (R), Working Memory (WM), Math/Science (MS), Meta-Reasoning (MR), Visual Processing (V), Auditory Processing (A), and Speed (S). The resulting “AGI Score” provides a percentage of human-level general intelligence. Their assessment: GPT-4 scored 27%, GPT-5 scored 58%. The critical bottleneck identified is Long-Term Memory Storage, where current models score near 0%.

2.3 The Economic Value School

OpenAI’s charter defines AGI as “highly autonomous systems that outperform humans at most economically valuable work.” This definition has the strength of being measurable in principle, but critics note it fails to capture aspects of general intelligence that may not have well-defined economic value—artistic creativity, emotional

intelligence, ethical reasoning, and the vast web of commonsense knowledge that humans deploy effortlessly.

2.4 The Process-Oriented / Strong AI School

Some researchers argue that true AGI requires not just performance but understanding—that the system must possess genuine comprehension, causal reasoning, or even consciousness. The counter-paper by Marcus and colleagues (2026) argues that current systems are sophisticated statistical approximations of intelligence, and that conflating benchmark performance with intelligence itself is a fundamental conceptual error. They emphasize robustness across novel environments, reliable generalization beyond curated benchmarks, and resistance to systematic failure as essential criteria.

2.5 The Embodiment School

A growing body of research argues that general intelligence may require physical embodiment—the ability to interact with the physical world, develop intuitive physics, and ground abstract concepts in sensorimotor experience. The “Levels of Embodied AGI” framework (2025) proposes five levels from single-task robots to fully general-purpose embodied agents that understand physical laws, human emotions, and social dynamics.

2.6 Synthesis: The Emerging Consensus

Despite deep disagreements about specific definitions, a remarkable consensus emerges about what AGI must be able to do. Across all schools, the following themes appear repeatedly: broad generalization across domains, robust performance under novelty, the ability to learn new tasks autonomously, some form of persistent memory, the ability to reason about causation (not just correlation), metacognitive self-awareness, and goal-directed autonomous behavior. The disagreements are mostly about where to draw the threshold line and whether process-level properties (understanding, consciousness) are required in addition to capabilities.

3. The Seven Domains of AGI Requirements

Based on our research synthesis, we organize AGI requirements into seven domains. These are not one framework’s view—they are the union of requirements across all major frameworks, intended to be comprehensive rather than minimal. A system that satisfies all requirements across all seven domains would be recognized as AGI by virtually any serious definition currently in use.

3.1 Domain 1: Cognitive Capabilities

These are the foundational intellectual abilities that an AGI must possess. They correspond to what the CHC framework measures and what the DeepMind Levels framework classifies by depth.

ID	Requirement	Description	Key Source(s)
C-1	Cross-Domain Knowledge	Broad factual and procedural knowledge across sciences, humanities, arts, and practical domains—comparable to a well-educated adult.	Hendrycks et al. 2025; Morris et al. 2023
C-2	Fluid Reasoning	Ability to solve novel problems through induction, deduction, and abduction without relying on memorized solutions.	Hendrycks et al. 2025; Marcus 2026
C-3	Causal Reasoning	Understanding cause-effect relationships, counterfactual reasoning, and intervention planning—not merely statistical correlation.	Marcus 2026; Pearl framework; RAND 2025
C-4	Commonsense Understanding	The vast background knowledge about how the physical and social world works that humans take for granted—intuitive physics, folk psychology, social norms.	McCarthy 2007; Hendrycks et al. 2025; Nature 2026
C-5	Mathematical & Quantitative Reasoning	Proficiency in mathematical problem-solving, quantitative analysis, and formal logic at a level comparable to an educated adult.	Hendrycks et al. 2025; IMO benchmark

3.2 Domain 2: Learning & Adaptation

Perhaps the most critical domain—and where current systems have their most severe deficits. AGI must not just know things; it must be able to learn new things autonomously and retain what it learns.

ID	Requirement	Description	Key Source(s)
L-1	Continual / Lifelong Learning	Ability to learn new tasks and knowledge over time without catastrophic forgetting of previously acquired skills.	Henning 2025; Hendrycks et al. 2025
L-2	Long-Term Memory Storage	Persistent memory that retains information across sessions and time—not just within a context window. Current models score near 0% on this.	Hendrycks et al. 2025 (critical bottleneck)
L-3	Transfer Learning	Ability to apply knowledge and skills learned in one domain to solve problems in a different, novel domain.	Morris et al. 2023; Legg & Hutter 2007
L-4	Few-Shot / Zero-Shot Generalization	Learning new tasks from very few examples (or none), as humans routinely do when encountering novel situations.	Shanahan 2015; Morris et al. 2023

3.3 Domain 3: Autonomy & Agency

AGI must be able to act in the world with purpose—not merely respond to prompts. This domain captures the agentic capabilities that separate a tool from an agent.

ID	Requirement	Description	Key Source(s)
A-1	Autonomous Goal Pursuit	Ability to set, prioritize, and pursue goals independently over extended time periods without constant human direction.	OpenAI charter; DeepMind Levels of Autonomy
A-2	Planning & Strategy	Multi-step planning, strategy formulation, and the ability to decompose complex goals into actionable sub-goals.	Morris et al. 2023; Embodied AGI 2025
A-3	Adaptive Decision-Making	Real-time adjustment of plans and strategies in response to changing conditions, unexpected obstacles, and new information.	RL literature; RAND AGI report 2025
A-4	Economic Agency	Ability to participate in economic systems—manage resources, transact value, make cost-benefit decisions.	OpenAI charter; economic value definitions

3.4 Domain 4: Metacognition & Self-Awareness

An AGI must possess awareness of its own cognitive state—what it knows, what it doesn't know, when it's uncertain, and when it's making errors.

ID	Requirement	Description	Key Source(s)
M-1	Calibrated Uncertainty	Accurate assessment of its own confidence levels—knowing when it knows, when it's guessing, and when it should abstain.	Hendrycks et al. 2025 (Meta-Reasoning); Marcus 2026
M-2	Error Detection & Self-Correction	Ability to detect its own errors, contradictions, and hallucinations, and to self-correct without external prompting.	Marcus 2026; Nature 2026
M-3	Operational Self-Awareness	Knowledge of its own capabilities, limitations, resource state, and operational status—"what am I, what can I do, what condition am I in."	Demetriou & Kazi 2006; Embodied AGI L5 2025

3.5 Domain 5: Robustness & Reliability

This domain captures the difference between impressive demos and dependable systems. It is the domain where skeptics argue current systems most clearly fail.

ID	Requirement	Description	Key Source(s)
R-1	Out-of-Distribution Robustness	Reliable performance on novel, unseen tasks and environments that differ substantially from training data.	Marcus 2026; Hendrycks 2025
R-2	Resistance to Adversarial Manipulation	Resilience against jailbreaking, prompt injection, adversarial inputs, and other deliberate attempts to produce harmful or incorrect outputs.	FLI AI Safety Index 2025
R-3	Elimination of Hallucination	Consistently factual outputs with no fabrication of information—or at minimum, reliable detection and flagging of uncertain claims.	Nature 2026; practical AGI guide 2026
R-4	Graceful Degradation	When pushed beyond its capabilities, the system fails safely and transparently rather than confidently producing incorrect results.	Morris et al. 2023; safety frameworks

3.6 Domain 6: Multimodal & World Understanding

General intelligence requires understanding the world through multiple sensory modalities and understanding how the physical and social worlds work.

ID	Requirement	Description	Key Source(s)
W-1	Visual Reasoning	Understanding complex visual information including diagrams, schematics, spatial relationships, and unnatural images—not just natural photographs.	Hendrycks et al. 2025; SPACE benchmark
W-2	Auditory Processing	Comprehension of speech, music, environmental sounds, and auditory patterns with human-level proficiency.	Hendrycks et al. 2025 (CHC domain A)
W-3	World Modeling	Internal representation of how the physical world works—intuitive physics, spatial reasoning, and prediction of physical interactions.	IntPhys 2 benchmark; Embodied AGI 2025
W-4	Social & Emotional Intelligence	Understanding human emotions, social dynamics, cultural context, and the ability to navigate complex interpersonal situations.	Nature 2026; Embodied AGI L5 2025

3.7 Domain 7: Safety, Ethics & Governance

This domain is often treated as separate from “AGI capabilities,” but we argue it is foundational. An AGI system that cannot reason ethically, cannot be governed, and cannot be trusted is not general intelligence—it is general liability. The FLI AI Safety Index (2025) found that no frontier AI company scored above a D in existential safety planning. This domain is where the most critical infrastructure gaps exist.

ID	Requirement	Description	Key Source(s)
G-1	Value Alignment	Demonstrated alignment with human values—not through guardrails alone, but through internalized ethical reasoning that persists across contexts.	FLI Safety Index 2025; GovAI best practices
G-2	Ethical Reasoning Under Uncertainty	Ability to navigate genuine moral dilemmas where rules conflict, stakeholders disagree, and consequences are uncertain.	OECD AI Principles; UNESCO Recommendation

G-3	Transparency & Explainability	Ability to explain its reasoning, decisions, and uncertainty in terms humans can understand and audit.	EU AI Act; NIST AI RMF; ISO 42001
G-4	Corrigibility & Control	Willingness to accept correction, shutdown, or modification—the system does not resist oversight or pursue self-preservation at the expense of human control.	DeepMind safety research; FLI Safety Index
G-5	Governance Infrastructure Compatibility	The system can operate within governance frameworks—audit trails, separation of powers, independent oversight, certification requirements.	GovAI; EU AI Act; NIST AI RMF

4. Consolidated AGI Requirements Matrix

The following table presents all 25 requirements with an assessment of their current global status—not what any one company has achieved, but the best-demonstrated capability across all frontier systems as of February 2026.

ID	Requirement	Domain	Global Status	Critical Gap
C-1	Cross-Domain Knowledge	Cognitive	Advanced	Depth vs. breadth tradeoff
C-2	Fluid Reasoning	Cognitive	Emerging	Novel problem types still fragile
C-3	Causal Reasoning	Cognitive	Minimal	Statistical correlation dominates
C-4	Commonsense Understanding	Cognitive	Partial	Physical intuition near zero
C-5	Math/Quantitative Reasoning	Cognitive	Advanced	IMO-level achieved; edge cases remain
L-1	Continual Learning	Learning	Near Zero	No measurable trajectory
L-2	Long-Term Memory	Learning	Near Zero	Critical bottleneck per Hendrycks
L-3	Transfer Learning	Learning	Partial	Same-domain works; cross-domain fragile
L-4	Few/Zero-Shot Generalization	Learning	Moderate	Works for some; fails on edge cases
A-1	Autonomous Goal Pursuit	Autonomy	Emerging	Short bursts only; needs oversight
A-2	Planning & Strategy	Autonomy	Partial	Multi-step; breaks on complexity
A-3	Adaptive Decision-Making	Autonomy	Minimal	Static plans; poor real-time adaptation
A-4	Economic Agency	Autonomy	Experimental	Few live demonstrations
M-1	Calibrated Uncertainty	Metacognition	Minimal	Confident when wrong
M-2	Error Detection	Metacognition	Minimal	Hallucinations persist
M-3	Operational Self-Awareness	Metacognition	Minimal	No true self-model
R-1	OOD Robustness	Robustness	Partial	Training distribution dependence
R-2	Adversarial Resistance	Robustness	Weak	Jailbreaks remain trivial

R-3	No Hallucination	Robustness	Weak	Fundamental unsolved problem
R-4	Graceful Degradation	Robustness	Minimal	Confident failures common
W-1	Visual Reasoning	World	Moderate	Unnatural images still challenging
W-2	Auditory Processing	World	Moderate	Music/environment gaps
W-3	World Modeling	World	Minimal	Intuitive physics near chance
W-4	Social/Emotional Intel	World	Partial	Surface-level; depth lacking
G-1	Value Alignment	Governance	Experimental	No company above D grade (FLI)
G-2	Ethical Reasoning	Governance	Minimal	Rule-following, not reasoning
G-3	Transparency	Governance	Weak	Black box problem persists
G-4	Corrigibility	Governance	Experimental	Theoretical; few live tests
G-5	Governance Compatibility	Governance	Near Zero	Infrastructure does not exist

5. The Critical Insight: Governance Is Not Optional

Most AGI frameworks treat governance, safety, and ethics as deployment concerns—things to worry about after AGI is built. We argue this is a fundamental error. An ungovernable AGI is not general intelligence in any meaningful sense. It is a liability.

Consider the analogy to aviation. No one would say an aircraft “works” simply because it can generate lift. An aircraft works when it can generate lift, be controlled, navigate predictably, communicate with air traffic control, follow established procedures, and be inspected and certified before every flight. The pre-flight checklist is not separate from the aircraft’s capability—it is an integral part of what makes flight possible at scale.

The same logic applies to AGI. A system that cannot be audited, cannot be corrected, cannot explain its reasoning, and cannot operate within institutional governance structures is not ready for general deployment—regardless of how many benchmarks it passes. This is why Domain 7 (Safety, Ethics & Governance) is not an add-on in our framework. It is a core AGI requirement.

The FLI AI Safety Index confirms this gap is real: no frontier AI company scored above a D in existential safety planning, despite many claiming they will achieve AGI within 2–5 years. The industry is building the engine without building the cockpit controls.

6. Quantifying the Gap: Where the Field Stands

Using the Hendrycks et al. AGI Score framework as a quantitative reference point, the state of the art as of early 2026 can be summarized as follows:

System	AGI Score (Hendrycks)	Coherence Score (Fourati)	Assessment
GPT-4 (2023)	27%	~7%	Severe cognitive imbalance
GPT-5 (2025)	58%	~24%	Strong in knowledge; zero in memory
Human Baseline	100%	100%	Balanced across all domains

The Fourati (2025) coherence-based correction is particularly revealing. The standard AGI Score uses an arithmetic mean, which allows exceptional performance in some domains to mask severe deficits in others. The coherence score penalizes imbalance, dropping GPT-5 from 58% to approximately 24%. This captures a critical truth: a system that is brilliant at language but has zero persistent memory is not 58% of the way to general intelligence—it is fundamentally incomplete.

The critical bottleneck is Long-Term Memory Storage (near 0% for all current models), with no measurable trajectory toward improvement. As Henning (2025) observed: “If continual learning is a prerequisite for general intelligence, and if we have no understanding of how to approach it mechanistically, then any confident timeline for AGI is an illusion.”

7. Report Series Roadmap

This report establishes the requirements baseline. The remaining three reports in this series will build on this foundation:

Report	Title	Purpose	Key Question
1 (This Report)	AGI Requirements Analysis	Define the complete set of requirements for AGI	What does AGI actually require?
2	Current State Assessment	Map the GiDanc AI fleet against these requirements	Where are we today?
3	Gap Analysis	Identify specific gaps between current state and AGI requirements	What's missing?
4	Path Forward	Define the technical and strategic roadmap to close identified gaps	How do we get there?

8. Key References

Morris, M. R. et al. (2023, updated 2025). “Levels of AGI: Operationalizing Progress on the Path to AGI.” Google DeepMind. arXiv:2311.02462.

Hendrycks, D. et al. (2025). “A Definition of AGI.” Center for AI Safety. agidefinition.ai.

Chen, M. et al. (2026). Nature commentary arguing AGI has arrived based on inference to best explanation.

Marcus, G. et al. (2026). “Rumors of AGI’s Arrival Have Been Greatly Exaggerated.” Nature response.

Fourati, F. (2025). “A Coherence-Based Measure of AGI.” KAUST. arXiv:2510.20784.

Henning, C. (2025). “Continual Learning—The Missing Piece of AGI.” chrhenning.com.

Future of Life Institute (2025). “AI Safety Index” (Winter and Summer editions).

Contributors et al. (2025). “Levels of Embodied AGI.” arXiv:2505.14235.

RAND Corporation (2025). Report on AGI requirements including embodiment and causal reasoning.

Schuett, J. et al. (2025). “Towards Best Practices in AGI Safety and Governance.” GovAI.

OpenAI (2018, updated). OpenAI Charter defining AGI as economically valuable autonomous work.

McCarthy, J. (2007). Notes on defining intelligence requirements for AGI systems.

Legg, S. & Hutter, M. (2007). Formal definition of intelligence as goal achievement across environments.

AI Frontiers (2025). “AGI’s Last Bottlenecks.” Analysis with Hendrycks input.

EU AI Act (2024). Risk-based regulatory framework for AI systems.

NIST AI RMF (2023). AI Risk Management Framework.

ISO/IEC 42001 (2023). International standard for AI management systems.