

An Engineering Definition of AGI: Complex-Project Closure as the Criterion

Version: V1.1

Author: Liu Ming

Affiliation: Independent Researcher, Beijing, China

ORCID: 0009-0001-2816-538X

Email: erran@zymobile.net

Date: September 3, 2026

Abstract

Artificial general intelligence (AGI) has long lacked a stable and operational definition. Existing definitions variously emphasize human-like behavior, subjective mind, broad cognitive ability, human-level performance, economic value, or social impact, yet they still struggle to answer an engineering question: when does AI cease to be a high-capability tool and become a system able to assume independent responsibility for complex production? This paper shifts the unit of judgment from local capabilities and individual tasks to complex-project closure and proposes a working definition of engineering AGI: given goals, resources, safety boundaries, and initial context, and where acceptance boundaries can be established independently, an AI system can, without continuous human direction, reliably complete most representative complex projects that previously required long-term collaboration by ordinary human teams, while producing verifiable deliverables. The definition primarily covers complex projects whose solution paths may be unknown but that can be closed by external goals and evaluation conditions; it does not require the system to consistently produce original contributions beyond the human frontier on open problems that lack a stable formulation or evaluation standard. The paper's structural-increment claim is that the key unit for judging AGI should move from local capability to continuity of responsibility in closable complex projects, and should be clearly distinguished from open-frontier innovation. Its directional-value claim is that this boundary makes AGI continuously measurable through project coverage, human intervention, delivery reliability, and independent acceptance evidence, providing a comparable, verifiable, and incrementally approachable system-level target for AI engineering, knowledge-work automation, and productivity assessment. The definition does not attempt to resolve philosophical questions such as consciousness or subjective experience.

Keywords: artificial general intelligence; engineering AGI; complex projects; closable projects; project closure; autonomous systems

1. Introduction

Artificial general intelligence has long been regarded as a central objective of AI development, yet there is still no shared basis for judging whether AGI is near. Turing translated the question of machine intelligence into an observable behavioral test [1]; Searle raised the philosophical objection that programs alone may not constitute understanding or mind [2]; and Legg and Hutter formalized general intelligence in terms of an agent's ability to achieve goals across a wide range of environments [3]. Human-level performance, substitution for economically valuable work, breadth and depth of capability, cognitive structure, and the effects of social transformation have since become major approaches to defining and measuring AGI [4-9].

Each approach answers an important question, but none fully closes the same engineering gap: why can an AI that performs well on many local tasks still be unable to take over a real complex project? Local tasks primarily test particular capabilities. Complex projects additionally require a system to preserve goals over time, manage task dependencies, respond to process changes, recover from errors, validate outcomes, and maintain what it delivers. As long as these functions still depend mainly

on continuous human direction, the AI remains a high-capability tool rather than a bearer of project responsibility.

This paper therefore proposes a working definition of AGI for engineering practice. It is not intended to replace every other definition of AGI, but to add a level of judgment closer to the process by which real productivity is created: whether AI can continuously organize its existing capabilities into the reliable delivery of complex projects.

2. A Working Definition of Engineering AGI

2.1 Core definition

This paper proposes:

Engineering artificial general intelligence is a system-level capability state in which an AI system, given goals, available resources, safety boundaries, and initial context, and where acceptance boundaries can be established independently, can, without continuous human direction, reliably complete most representative complex projects that previously required long-term collaboration by ordinary human teams, while producing verifiable deliverables.

In simplified form:

Engineering AGI is the capacity of AI to complete, without ongoing human operation, complex projects that previously required long-term collaboration by ordinary human teams.

The simplified statement is suitable for public communication; the formal definition is intended for research and evaluation. The central question is not whether AI resembles a human, but whether AI can reliably assume the principal cognitive work, process governance, and delivery judgments required to move a project forward.

2.2 Key concepts

An **AI system** is not a foundation model in isolation, but an operational whole comprising a model together with memory, tools, an execution environment, and necessary governance mechanisms. The same model can exhibit substantially different project capabilities under different system configurations. Engineering AGI is therefore a system-level property.

A **complex project** is not merely a long task. It is a unit of work containing multilayer dependencies, intermediate goals, process feedback, and requirements for an integrated deliverable. It usually requires sustained planning, execution, verification, and correction under incomplete information and changing conditions. For ongoing operations without a natural endpoint, a bounded operating period with a clear objective, time horizon, and acceptance conditions can serve as the unit of project closure.

A **closable project** does not imply that the solution path is known in advance, or that every local step has a unique correct answer. It means that the project's objectives, principal constraints, and acceptance boundaries can be established independently before the system produces its final solution, or can stabilize during the project. Software development, experimental replication, product implementation, and many analytical projects may all involve incomplete information, local exploration, and design innovation. They remain within the closable scope of this paper so long as the final deliverable can be accepted against conditions that are relatively independent of the proposed solution.

An **ordinary human team** is the reference baseline for project difficulty. It means a professional team with the relevant occupational competence, using conventional tools and organizational processes-not

an average individual and not a special team composed of top experts. A concrete evaluation should state the human-team baseline used, including its typical resources, time horizon, and capability range.

Without continuous human direction does not mean that no human is involved. Humans may still set goals, allocate resources, define safety boundaries, and retain final authority. The AI, however, no longer depends on humans to continuously decompose tasks, correct direction, restore context, repair the execution process, or perform routine acceptance work on its behalf.

Reliable completion does not mean a single success or a formally complete-looking output. It requires repeatable delivery capability across similar projects, results that can undergo independent acceptance, and no appearance of success obtained by narrowing requirements, concealing failures, or transferring critical work to humans.

2.3 Project closure

In this paper, project closure means that a system can begin from a given objective, continuously organize the principal cognitive and governance activities required by the project, produce an acceptable result, and preserve delivery capability when errors or reasonable changes occur. Closure concerns continuity of responsibility, not whether the AI personally performs every local action.

An AI may therefore carry out some real-world operations through software interfaces, automated equipment, robots, professional services, or instructed humans. External humans may provide authorization, physical execution, or clearly delimited service interfaces, but they may not be used to supply project planning, critical judgment, process correction, or routine quality control that the system itself lacks. So long as the AI performs the core judgment, coordination, and quality control required by the project, such external interfaces do not change the nature of project closure.

Project closure permits a system to generate new solutions, new representations, and local structural innovations during execution. Engineering AGI differs from higher-level open intelligence not because it generates no new structure, but because those structures still serve a project objective that can be closed externally and can be judged by delivery conditions independent of the final solution.

3. Why Complex-Project Closure Is the Key Unit of Judgment

3.1 Local intelligence does not directly imply a productivity transition

Current AI can already write code, generate documents, analyze materials, assist design, and complete many independent tasks. In most real settings, however, humans still formulate specific problems, decompose work, coordinate dependencies, judge results, correct errors, and assume maintenance responsibility. AI has increased local efficiency without yet producing a general change in the structure of project responsibility.

The transition to engineering AGI occurs when AI is no longer merely a capability module that is invoked, but begins to assume responsibility for project progress. Humans move from continuous drivers to goal setters, boundary designers, risk auditors, and final authorizers.

This change connects more directly to social productivity than local-task performance does. Important outputs such as software systems, research programs, manufacturing processes, educational platforms, and organizational operations all fundamentally depend on complex-project capability. Only when the costs of project execution, coordination, and maintenance fall substantially can many projects that were previously infeasible because of organizational and trial-and-error costs become viable.

This paper uses a team rather than an individual human as the reference, not because AGI must equal the sum of multiple people's cognitive abilities, but because the natural unit of responsibility in modern complex production is usually a team. The engineering definition measures whether an AI

system can take over that unit of productive responsibility, not whether it reaches a particular human level on an individual cognitive test.

3.2 Project capability is not a simple sum of local capabilities

Complex projects are path-dependent. Errors in early requirements analysis, technical choices, or resource allocation may be amplified in later stages; a locally correct result can also undermine the coherence of the project as a whole. Project closure therefore requires, at minimum, cross-stage goal retention and global coordination.

These capabilities may be supplied jointly by external memory, audit procedures, and tool systems; they need not all reside within one model. Regardless of implementation, the final judgment should concern whether the system can continuously produce reliable deliverables, not how many isolated tasks it scores highly on.

Recent work that calibrates the scale of tasks AI can complete reliably by the time required for a human baseline has begun to move evaluation from static questions toward longer-horizon work [10]. Benchmarks such as PaperBench go further by requiring an AI agent to understand a research paper from scratch, develop code, and run experiments, approaching an end-to-end form of evaluation for complex work [11]. These directions show that task duration, reliability, and end-to-end execution have independent value. Real complex projects, however, also contain multilayer dependencies, process changes, maintenance responsibility, and external collaboration, so they cannot be reduced to task length or to end-to-end tasks in a single domain.

4. Relationship to Existing Definitions of AGI

Existing AGI definitions select different units of judgment. They do not simply compete with the present definition; rather, they describe different layers of artificial intelligence.

Approach	Primary focus	Relevance and limitation for engineering judgment
Human-like behavior	Whether a machine behaves like a human [1]	Provides an external behavioral criterion, but human-like behavior is not equivalent to project-delivery capability
Mind and understanding	Whether a machine genuinely understands or possesses a mind [2]	Addresses important philosophical questions, but is difficult to use as an engineering acceptance criterion
Formal general intelligence	An agent's ability to achieve goals across broad environments [3]	Offers a more general theoretical abstraction, but remains distant from evidence about real projects
Human-level and cognitive performance	Breadth, depth, and human-relative level of capabilities [4,6-8]	Can describe capability structure, but cannot by itself establish long-term continuity of responsibility
Economic work	Whether AI surpasses humans at most economically valuable work [5]	Connects AGI to productivity, but operates at a relatively macro level
Social impact	Whether AI causes an Industrial Revolution-scale social transformation [9]	Describes consequences well, but can usually be confirmed only after the fact

Cognitive measurement can indicate which foundational capabilities a system has; task evaluation can indicate what a system can complete under given conditions; project-closure evaluation tests whether those capabilities can be continuously organized into reliable delivery. The three can provide evidence for one another, but they cannot substitute for one another.

The engineering definition is closest to economic-work definitions. A system that can broadly take over complex projects is likely to have major economic effects, but economic outcomes also depend on cost, institutions, regulation, and deployment speed. The engineering definition moves the point of judgment upstream to system capability itself, so confirmation need not wait until macroeconomic effects have already occurred.

5. Boundaries of the Definition

To avoid overextending the definition, this section states its principal boundaries.

5.1 An engineering working definition and the boundary of human responsibility

This definition is suitable for discussing AI engineering, knowledge-work automation, complex-project delivery, and substitution within organizational production. It does not directly answer whether a machine has consciousness, subjective experience, or genuine understanding. A system may satisfy this definition while its mental status remains disputed; conversely, even if some form of machine consciousness existed, it would not necessarily possess complex-project capability.

"Assuming project responsibility" is used functionally here: the AI can preserve goals, organize execution, surface problems, drive recovery, and produce acceptable results. Humans remain responsible for goal setting, resource authorization, safety boundaries, high-risk ethical judgments, and ultimate social accountability. Engineering AGI describes a transfer of responsibility for project execution and governance; it does not imply that AI automatically becomes a legal or moral agent.

5.2 Engineering AGI does not require full embodiment

Real complex projects already depend on teams, tools, and execution interfaces; no single person performs every physical operation. Whether an AI has a human-like body should not be a universal prerequisite for AGI across all domains. Robots and specialized automated equipment will expand the range of projects a system can undertake, but they are not necessary for the definition to hold.

5.3 Project autonomy in a single domain is not general intelligence

A system may achieve strong project closure in a single domain such as software development, drug screening, or manufacturing scheduling, but this establishes only domain capability. The "general" in AGI requires sufficient breadth of coverage and cannot be supported by a few specially designed projects, a single vertical system, or successful cases selected after the fact.

"Most representative complex projects" is empirically testable only relative to a preregistered project set or a representative sampling distribution. Any empirical claim that a system has reached engineering AGI should specify the project scope, sampling principles, and domain coverage, rather than treating post hoc selected successes as evidence of general capability.

5.4 Reliable completion, verifiable delivery, and the closability boundary

No engineering system is infallible. Reliable completion means that, under an explicit project scope, resource conditions, acceptance criteria, and risk level, the system achieves repeatable reliability across similar projects at a level sufficient to assume delivery responsibility. It does not require every project to succeed. Different domains may set different thresholds, and high-risk domains must impose stricter safety requirements.

"Verifiable delivery" also does not require every project to have a unique objective truth. In domains such as software and manufacturing, outcomes may be tested against explicit acceptance criteria. For work such as experimental replication, product implementation, option comparison, and strategic analysis, where paths are open but results can be closed, external acceptance should use criteria set in advance or established independently. The core requirement is that the system cannot be the sole source of the claim that its deliverable succeeded.

If the central difficulty of a problem is to discover what is worth studying, reformulate the problem, introduce new evaluative dimensions, or decide whether a new structure without an established standard is worth retaining, then the problem has not yet been sufficiently closed. Engineering AGI may participate in such work and may occasionally make important contributions, but the present definition does not require the system to reliably exceed the human cognitive frontier on such problems.

5.5 This paper does not prescribe a single evaluation standard

Project structure, risk, and acceptance methods differ greatly across domains, so it would be premature to prescribe a unified metric or fixed threshold. Project coverage, scale of task dependencies, degree of human intervention, delivery reliability, failure recovery, continuity of maintenance, completion time, and resource cost are all possible measurement dimensions. Concrete evaluation systems must be developed through further domain-specific research.

The absence of one universal metric does not mean that evaluation can omit operational conditions. Any concrete comparison or claim of engineering AGI should state, at minimum: the project set or sampling distribution, the ordinary-human-team baseline, allowed resources and time, rules for human intervention, acceptance conditions, and success criteria. Only when these conditions are explicit do "most," "representative," "reliably complete," and "without ongoing human operation" acquire comparable empirical meanings.

5.6 The boundary between engineering AGI and ASI

Some existing frameworks place artificial superintelligence (ASI) at the highest level of broad human-relative performance, for example by requiring a system to exceed every human across a wide range of tasks [6]. Research on open-ended intelligence further argues that continually producing outcomes that are novel and learnable to human observers is an important condition for artificial superintelligence [12]. This paper does not fully define ASI, but it must make the upper boundary of engineering AGI clear: engineering AGI requires broad autonomous delivery of closable complex projects, not the sustained generation, on genuinely open problems, of directions, abstractions, or evaluative frameworks beyond the human frontier.

Accordingly, the central change from current AI to engineering AGI is that responsibility for complex-project execution and governance moves from continuous human direction to autonomous AI closure. Moving from engineering AGI to ASI requires a further discussion of direction discovery, structure generation, and value judgment in open problems. These two changes cannot be summarized by the same set of local-task scores.

6. How an Engineering Definition Makes AGI an Approachable Target

The value of an engineering definition is not that it immediately supplies one universal score, but that it turns AGI from a vague imagined state into an observable direction of capability. Researchers need not only debate whether a system "resembles a human" or already possesses some form of general intelligence. They can ask: what scale, duration, and range of projects can it assume? Is the project's dependence on human direction declining? Are delivery reliability, failure recovery, and long-term maintenance improving?

These questions can generate continuous evidence. Even before any system reaches the final threshold, different systems can be compared in project scope, complexity, and degree of autonomy once the project distribution, resource conditions, and human baseline are explicit. Improvements in individual capabilities can thereby be understood within a fuller production structure, and AGI no longer has to be treated as a binary event awaiting announcement.

This view also explains an apparent paradox in current AI development: rapid gains in local capability have not translated to the same degree into organizational change or gains in social productivity. The missing link is not one more local skill, but the system-level capacity to continuously organize multiple capabilities into project closure.

7. Discussion: From Capability Evaluation to Continuity of Responsibility

Existing capability benchmarks remain important components of engineering-AGI evaluation, but they should sit within a higher-level evidence structure. Local capability tests show which capabilities a system possesses; task evaluations show what it can complete under given conditions; project evaluations test whether those capabilities can be organized across stages into maintainable deliverables.

Complex-project closure therefore does not replace existing evaluations with a single new benchmark; it changes the final unit of judgment. As systems improve, evaluation should gradually shift from one-off capability demonstrations toward project scope, degree of autonomy, failure recovery, long-term maintenance, and continuity of responsibility. The defining sign of engineering AGI is not merely how intelligent a system appears at a given moment, but whether it can reliably take over a unit of productive responsibility that previously depended on sustained collaboration by a human team.

8. Conclusion

This paper proposes complex-project closure as the engineering criterion for AGI: given goals, resources, safety boundaries, and initial context, and where acceptance boundaries can be established independently, an AI system can, without continuous human direction, reliably complete most representative complex projects that previously required long-term collaboration by ordinary human teams, while producing verifiable deliverables.

The definition is not intended to replace cognitive, formal, or social-impact definitions of AGI, but to add a layer of project responsibility that they do not fully cover. It extends the object of judgment from an isolated model to a complete system, expands the unit of progress from local tasks to complex projects, and makes the framework more operational and comparable by requiring explicit project scope, human baseline, human intervention, resource conditions, and acceptance methods.

Only when AI can continuously organize broad local capabilities into reliable, autonomous, and maintainable delivery of complex projects does it move from a high-capability tool to a genuine system of project responsibility. Taking this transition as the unit of judgment makes engineering AGI a comparable, verifiable, and incrementally approachable system-level engineering target. At the same time, limiting this target to complex projects that can be closed by external conditions avoids conflating project-execution autonomy with open-frontier innovation and provides a clear interface for further analysis of the engineering boundary between AGI and ASI.

References

- [1] TURING A M. Computing machinery and intelligence[J]. *Mind*, 1950, 59(236): 433-460. DOI: [10.1093/mind/LIX.236.433](https://doi.org/10.1093/mind/LIX.236.433).
- [2] SEARLE J R. Minds, brains, and programs[J]. *Behavioral and Brain Sciences*, 1980, 3(3): 417-457. DOI: [10.1017/S0140525X00005756](https://doi.org/10.1017/S0140525X00005756).
- [3] LEGG S, HUTTER M. Universal intelligence: A definition of machine intelligence[J]. *Minds and Machines*, 2007, 17(4): 391-444. DOI: [10.1007/s11023-007-9079-x](https://doi.org/10.1007/s11023-007-9079-x).
- [4] GRACE K, SALVATIER J, DAFOE A, et al. When will AI exceed human performance? Evidence from AI experts[J]. *Journal of Artificial Intelligence Research*, 2018, 62: 729-754. DOI: [10.1613/jair.1.11222](https://doi.org/10.1613/jair.1.11222).
- [5] OPENAI. OpenAI Charter[EB/OL]. 2018. <https://openai.com/charter/>.
- [6] MORRIS M R, SOHL-DICKSTEIN J, FIEDEL N, et al. Position: Levels of AGI for operationalizing progress on the path to AGI[C]//*Proceedings of the 41st International Conference on Machine Learning*. PMLR, 2024, 235: 36308-36321. arXiv: [2311.02462](https://arxiv.org/abs/2311.02462).
- [7] HENDRYCKS D, SONG D, SZEGEDY C, et al. A definition of AGI[EB/OL]. arXiv, 2025. arXiv: [2510.18212](https://arxiv.org/abs/2510.18212).

- [8] BURNELL R, YAMAMORI Y, FIRAT O, et al. Measuring progress toward AGI: A cognitive framework[EB/OL]. arXiv, 2026. arXiv: [2605.28405](#).
- [9] KARNOFSKY H. Some background on our views regarding advanced artificial intelligence[EB/OL]. Open Philanthropy, 2016. <https://coefficientgiving.org/research/some-background-on-our-views-regarding-advanced-artificial-intelligence/>.
- [10] KWA T, WEST B, BECKER J, et al. Measuring AI ability to complete long tasks[EB/OL]. arXiv, 2025, revised 2026. arXiv: [2503.14499](#).
- [11] STARACE G, JAFFE O, SHERBURN D, et al. PaperBench: Evaluating AI's ability to replicate AI research[EB/OL]. arXiv, 2025. arXiv: [2504.01848](#).
- [12] HUGHES E, DENNIS M, PARKER-HOLDER J, et al. Position: Open-endedness is essential for artificial superhuman intelligence[C]//Proceedings of the 41st International Conference on Machine Learning. PMLR, 2024, 235: 20597-20616. arXiv: [2406.04268](#).