

KSC2: A Progressive Structural Framework for ASI — Structural Openness, Structural Autonomy, and the Structural Difference between AGI and ASI

Liu Ming

Version 1.2 • August 2026

Abstract

Artificial General Intelligence (AGI) and Artificial Superintelligence (ASI) are usually characterized in terms of capability coverage and performance relative to humans. Such descriptions, however, do not directly capture whether AI can autonomously form more effective problem representations, abstraction schemes, and exploration directions when existing problem structures are inadequate for solving a task. This paper asks a further question: **before complete ASI is achieved, can progress in this capability be continuously observed and compared at the level of task structure?**

Starting from the Knowledge–Structure–Constraint Framework (KSC), this paper introduces **Structural Openness**, which describes, under a fixed task protocol, how much of the high-level problem representations, abstraction schemes, key variables, relational structures, and exploration directions required to complete a task is not externally specified before task execution begins and must instead be formed during problem solving. Structural openness is therefore a task-side coordinate of structural demand rather than a measure of solver capability.

The paper further uses **Structural Autonomy** to describe AI’s ability to autonomously form, evaluate, and revise the high-level structures required to complete a task in the absence of critical human S input during task execution. Structural autonomy is a system-side coordinate of structural capability. On this basis, structural progress toward ASI can be observed as AI reliably takes on tasks with increasing structural openness, steadily reduces its dependence on critical human S input while maintaining effective structure generation, and extends this capability across an increasingly broad range of domains.

This framework does not replace classic capability-level definitions of ASI. It describes one important capability dimension required for complete ASI and cannot, by itself, measure how far a system remains from complete ASI. The paper further argues that the difference between AGI and ASI should not be understood as “AGI has no structural innovation, whereas ASI does,” but as a difference in the **required breadth of coverage and capability ceiling of structural autonomy**: the central goal of AGI is broad autonomous coverage of intellectual work that ordinary humans can complete through sustained collaboration; complete ASI additionally requires that, across a broad range of high-structural-openness tasks, AI autonomously enter areas of capability that have historically depended mainly on the most capable humans for problem reframing, abstraction discovery, and original structure generation, and reach or exceed the highest human level. AGI and ASI can therefore be understood as two capability frontiers that may advance in parallel and interact with one another rather than as stages that must occur in a strict sequence.

Keywords: KSC; Artificial General Intelligence; Artificial Superintelligence; Structural Openness; Structural Autonomy; Structure Generation; Open-Ended Intelligence; Human–AI Collaboration

1. Introduction

Generative AI is rapidly expanding its capabilities in knowledge use, reasoning, planning, code generation, and complex task execution. As models and agent systems continue to develop, many specific generative tasks that previously required substantial human participation are gradually becoming machine-executable. Yet the important role of humans has not simply disappeared. Instead, it is becoming increasingly concentrated at a small number of high-level decision points that determine how a problem is understood and what direction a task should take.

The KSC framework interprets intelligent activity in complex goal-oriented tasks in terms of three ordered functional levels: a knowledge foundation (K), structure generation (S), and constraint maintenance (C). K provides the information and state available to the task; S forms problem representations, relationships, solutions, and action paths from that information; C ensures that established goals, premises, boundaries, and acceptance criteria continue to govern generation and state transitions [1].

In tasks where goals and acceptance conditions are relatively clear, AI can take on much of K and S, while the primary human role is often concentrated in constraint maintenance. However, when it remains uncertain how the problem itself should be understood, which variables matter, what abstractions are effective, or which direction should be explored next, high-level human guidance no longer merely constrains existing generation. Redefining the problem, discovering new relationships, or proposing new directions of exploration directly generates new S.

This distinction is especially relevant to how we understand capability progress in relation to AGI and ASI.

There is still no single universally accepted definition of AGI. Morris et al. developed a graded framework for AGI based on performance level and generality, allowing AGI to be described as a process of continuous progress rather than a single binary milestone [2]. The OpenAI Charter, meanwhile, described AGI as a highly autonomous system that outperforms humans at most economically valuable work [3]. Burnell et al. further proposed a framework for measuring progress toward AGI based on decomposed cognitive capabilities [7].

Classic definitions of ASI place greater emphasis on performance relative to the upper end of human capability. Bostrom described superintelligence as intelligence that is substantially more capable than the best human brains in virtually every important domain, explicitly including scientific creativity, general wisdom, and social skills [4,5]. Hughes et al. further proposed that open-endedness—the sustained production of outcomes that are novel and learnable relative to an observer—is an important condition for artificial superintelligence [6].

These lines of research primarily ask what level of capability would qualify a system as AGI or ASI, and how overall progress toward those capability states may be described. This paper focuses on a more specific question:

When existing problem structures are insufficient for solving a task, can AI autonomously discover, form, and validate more effective high-level structures? Can this capability serve as an independent dimension for observing progress toward ASI?

2. Structural Openness: When the Problem Space Itself Becomes Part of the Problem

2.1 Complexity Does Not Equal Openness

A complex task does not necessarily have high structural openness.

A large software system may contain millions of lines of code, many modules, and long-term collaboration, while still having numerous local tasks with clear goals, mature technical frameworks, observable states, and testable outcomes. Although the overall task is large in scale, the problem-solving process can still be decomposed and progressively constrained into relatively stable problem spaces.

Conversely, a very short problem statement may have high structural openness. For example:

- How should an anomalous phenomenon actually be understood?
- Has a research program that has stalled for years been asking the wrong question from the beginning?
- Do two apparently independent phenomena share the same underlying mechanism?
- Are the variables used by the current theory partitioned incorrectly?
- What problem should a new product actually solve?

The main difficulty in these cases is not that there are too many steps, but that **the high-level structures required to complete the task have not been specified in advance.**

2.2 Structural Openness Depends on the Task Protocol

This paper defines **Structural Openness** as:

Given an initial problem representation, externally supplied knowledge and tools, feedback conditions, and an evaluation protocol, structural openness is the extent to which the effective problem representations, abstraction schemes, key variables, relational structures, and exploration directions required to complete a task are not specified before task execution begins and therefore must be formed during the solving process.

This relation can be conceptually represented as:

$$O(T \mid R_0, K_0, E)$$

where:

- T denotes the target task;
- R_0 denotes the externally supplied problem representation and structure at the start of the task;
- K_0 denotes the knowledge, tools, and initial state externally provided by the task protocol;
- E denotes the feedback, validation, and evaluation environment;
- O denotes the remaining high-level structural requirements not specified in advance under those conditions.

Here, K_0 does not represent the full body of training knowledge internalized in the parameters of the evaluated system. Knowledge accumulated within the model is part of the system's capability. If a high-level structure is not externally provided as part of the current task protocol, but the system can retrieve, recover, reconstruct, or regenerate that structure using its own capabilities, the difference should be reflected in the system's structural autonomy rather than changing the task's structural openness.

Therefore, when comparing different models or systems, T , R_0 , K_0 , and E should be held fixed as much as possible. Different systems face the same task conditions; the distinction lies in whether they can autonomously form the structures that have not been given.

Structural openness is therefore neither an absolute property of a task independent of the conditions under which it is solved, nor a function of solver capability. It describes **how much critical high-level structure still has to be formed during problem solving under a clearly specified task protocol**.

It is also not equivalent to whether a task has a unique correct answer, nor is it synonymous with task value, scale, or difficulty in the general sense.

2.3 Structural Openness Is Currently Better Suited to Graded Comparison

This paper does not currently claim that O can be stably computed as a unified numerical quantity.

The structural requirements not specified in advance may involve multiple dimensions at the same time:

- whether the correct problem representation has already been provided;
- whether important variables have already been provided;
- whether candidate relationships have already been provided;
- whether the system must choose among multiple candidate structures;
- whether it must actively overturn an existing structure;
- whether it must change the level of abstraction;
- whether it must propose exploration directions not originally present in the task protocol.

These changes do not naturally collapse into a single scalar. At the current stage, it is therefore more appropriate to compare structural openness in graded form under controlled task conditions and to establish locally comparable relations.

For example, for the same underlying problem, one could construct:

The correct problem representation is directly provided

→ **Multiple candidate representations are provided, but the correct structure is not identified**

→ **Only the phenomenon, materials, and objective are provided**

→ **The initial problem description itself contains a misleading structure that must be actively overturned.**

Within such a controlled task family, successive variants require the system to assume more responsibility for problem formation, structure selection, and structural revision, and can therefore be considered to impose higher structural-openness requirements.

Accordingly, when this paper refers to “higher structural openness,” it primarily means an increase in structural requirements under such controlled conditions, rather than implying the existence of a unified numerical scale across all tasks.

2.4 From Searching for Answers to Constructing the Problem Space

In tasks with relatively low structural openness, the system usually already has a relatively stable problem representation, a mature body of methods, clear goal boundaries, and feedback mechanisms capable of eliminating incorrect paths.

The central question is then:

What path within the given problem space can produce an effective answer?

As structural openness increases, the current problem representation itself may be wrong, important variables may not yet have been discovered, candidate relationships may need to be formed anew, and even the evaluation dimensions may change during exploration.

The question then becomes:

What kind of problem space should be constructed in order to search for an answer?

This second question is itself part of structure generation S.

Structural openness therefore describes the extent to which task solving shifts from “primarily solving within an existing problem space” toward “forming or revising the problem space itself during solving.” It is not intended to divide engineering, scientific research, or creative activity into fixed categories.

3. From Constraint Collaboration to Structural Collaboration

3.1 Not All High-Level Human Guidance Belongs to C

Within KSC, the core role of constraint maintenance C is to ensure that already-formed goals, premises, boundaries, and acceptance criteria continue to govern subsequent generation [1].

For example:

The current solution violates the established performance requirement.

This primarily falls under C.

Likewise:

The current result has failed the predefined test and must be rolled back.

This also primarily falls under C.

But if a human instead says:

This may not be a performance problem at all; it may be a state-consistency problem.

the system receives a new problem representation.

If the human proposes:

Two phenomena that were previously handled separately may share the same underlying mechanism.

the system receives a new relational structure.

If the human judges:

Continuing to optimize along the current direction is meaningless; the problem needs to be redefined at a higher level of abstraction.

then the subsequent search space itself changes.

According to the functional boundaries of KSC, these actions directly participate in structure generation (S) and should not all be subsumed under C.

3.2 Human-AI S Collaboration

One of the distinctive capabilities of current AI is its ability to expand an existing structure at relatively low marginal cost.

Once given a valuable problem representation, hypothesis, or exploration direction, a model can rapidly connect relevant knowledge, elaborate lines of reasoning, search for

counterexamples, compare structures, conduct local validation, and derive implementations. As a result, humans no longer need to carry out all of the cognitive work that follows a single critical abstraction.

A typical form of human–AI collaboration on open problems can be represented as:

Critical structure generation → large-scale AI expansion → new knowledge and candidate structures → higher-level structure generation → renewed expansion.

The critical structure in this process may come from either a human or AI.

In many current systems, what may be truly scarce is not total reasoning volume, but a small number of high-value S inputs that substantially change the subsequent problem space. Recent research on autonomous science likewise suggests that automating the execution of scientific workflows and autonomously forming, challenging, and revising the high-level structures that determine research direction remain distinct levels of capability [8–10].

Human–AI collaboration is therefore undergoing an important shift:

from humans primarily constraining machine generation within existing structures, toward humans and machines jointly forming new problem structures.

4. Structural Autonomy and Structural Progress Toward ASI

4.1 Structural Autonomy

This paper uses **Structural Autonomy** to describe a system’s capability when the high-level structures required for a task have not been specified in advance:

Structural autonomy is the ability of a system to autonomously form, evaluate, and revise the high-level structures required to complete a task in the absence of critical human S input during task execution.

Here, “formation” may include recovering, reconstructing, or combining effective structures from existing knowledge, as well as generating new problem representations, variables, relationships, and explanatory frameworks.

Structural autonomy is distinct from how long an agent can operate autonomously.

One system may work independently for days while remaining trapped within the wrong problem representation. Another system may discover through a brief reasoning process that the original problem definition is wrong and form a more effective abstraction.

Structural autonomy is therefore not concerned with how long AI can continue executing by itself, but rather with whether:

when the current structure is insufficient for solving the problem, AI can form a more effective structure by itself and validate that change with external evidence.

This paper considers only the system’ s ability to form structures during the current task execution process. It does not require those structures to be historically novel in human knowledge.

If a structure already exists in human knowledge, and may even have appeared in the model’ s training material, but is not directly provided as part of the current task protocol, and the system can autonomously recover, combine, or reconstruct it and thereby alter the subsequent problem space, this capability can still count toward structural autonomy.

Demonstrating that a system has produced an original structure that did not previously exist in human knowledge would require much stricter verification of novelty and knowledge provenance. That is not a necessary condition under the definition used in this paper.

4.2 Critical Human S Input

The structural autonomy discussed here does not mean that AI can generate all cognitive structures independently of the accumulated achievements of human civilization.

Model training data, tool design, language, scientific theories, task definitions, and evaluation protocols all contain substantial structure accumulated through human history. The paper therefore examines a more limited and operational question:

Given a model, tools, initial task materials, and an evaluation environment, which critical human S inputs are still required to complete the current task?

“Critical” here has a counterfactual meaning and cannot be measured simply by counting prompts. One problem reframing may matter more than dozens of local hints; conversely, multiple hints that are not individually decisive may collectively change the outcome.

Accordingly, one can remove, replace, or delay a particular human structural input and observe whether the task success rate falls significantly, whether the system remains persistently trapped in an incorrect problem representation, or whether an effective solution path that was previously reachable no longer appears.

It is also necessary to distinguish validation from the provision of new structure.

If a human merely returns “pass/fail” according to a predefined evaluation protocol, or supplies the result of an external experiment, that primarily provides validation or maintains constraints. But if the evaluator additionally tells the system why the problem representation is wrong, which variable should be introduced, which abstraction level should change, or which direction should be explored, the validation process has also supplied new S.

Therefore, “AI proposes the structure and humans only validate it” can be considered a state of lower human S dependence only when the validation process does not additionally provide critical structural direction.

4.3 How to Observe Structural Progress Toward ASI

Classic definitions of ASI primarily answer:

What kind of system can be called superintelligent?

This paper poses a complementary question:

Before that endpoint is reached, what changes can count as progress toward the higher-order structural capabilities required by ASI?

This paper proposes:

Structural progress toward ASI can manifest as AI successfully handles an increasingly broad range of tasks with increasingly high structural openness, while steadily reducing dependence on critical human S inputs and progressively acquiring the ability to autonomously form, evaluate, and revise high-value structures.

Structural autonomy is not equivalent to complete ASI. Even if a system demonstrates very high structural autonomy in a narrow domain, that alone does not make it complete ASI; classic definitions of ASI still require extremely high capability levels across broad and important domains.

The claim of this paper is therefore limited: if a system is ultimately to reach or surpass the highest human levels across broad and important intellectual domains, it will likely need to depend increasingly less on humans to provide decisive task-specific problem reframing, abstraction discovery, and high-level direction selection.

4.4 Four Observational Dimensions

At the current stage, this paper does not compress structural progress toward ASI into a single numerical score. Instead, it proposes four dimensions that can be progressively operationalized through experiments.

First, Structural Openness of the Task

With R_0 , K_0 , and E fixed or recorded, observe the level of structural openness that AI can handle reliably.

Second, Dependence on Critical Human S Input

For the same task, one can observe a progression from:

Humans generate S; AI expands it

toward:

AI proposes candidate S; humans select

→ AI proposes and filters S; humans validate

→ AI autonomously proposes, filters, validates, and replaces S.

This progression can directly reflect increasing structural autonomy.

Third, Effectiveness of Structure Generation

Generating more directions does not by itself demonstrate higher structural autonomy.

An effective structure must show value through external evidence—for example, explaining more facts, substantially compressing the problem, producing new testable predictions, opening previously unavailable valid search directions, or improving actual outcomes in engineering, scientific research, or decision-making.

Fourth, Breadth of Capability Coverage

Superhuman structure generation in a local domain does not equal complete ASI. It is also necessary to observe how many different tasks and domains this capability can cover, and whether it can gradually expand across broad areas of human intellectual activity.

Together, these four dimensions can describe progress in the structural capabilities relevant to ASI without immediately collapsing them into a single score.

5. Relation to Existing Research on AGI, ASI, and Open-Ended Intelligence

5.1 Gradual Descriptions of AGI

Morris et al. and Burnell et al. have already discussed gradual measurement of AGI from the perspectives of capability level, generality, and cognitive capability profiles [2,7]. Therefore, this paper does not claim that the mere idea that “AGI can be described progressively” is itself new.

This paper instead asks:

Can a continuous coordinate related to task structure and structural autonomy be added to the process of approaching ASI?

5.2 Classic Definitions of ASI Emphasize the Final Capability Level

Bostrom’s classic definition of superintelligence uses the best humans as a comparison baseline and requires substantially higher capability across virtually all important intellectual domains, explicitly including scientific creativity, general wisdom, and social skills [4,5].

Accordingly, this paper does not claim that prior ASI research neglects creativity, scientific discovery, or higher-order cognition.

The different question raised here is:

Why do these higher-order capabilities impose different requirements on problem structure, and can those requirements be observed continuously before the ASI endpoint is actually reached?

Many complex forms of work can be completed within mature human cognitive structures. But in the highest levels of scientific, theoretical, and creative activity, some decisive progress occurs precisely when old structures are no longer sufficient.

From the KSC perspective, this means that AI must not only perform better within a given problem structure, but must also improve its autonomous capability when **the structure itself becomes an object of problem solving**.

5.3 Structural Openness Is Not the Same as Open-Endedness

Hughes et al. proposed that open-endedness is an important attribute of artificial superintelligence and formalized it in terms of the sustained production of outcomes that are novel and learnable relative to an observer [6].

The concept of **structural openness** in this paper is adjacent to, but distinct from, open-endedness because the two concepts characterize different objects.

Open-endedness primarily asks:

Can a system or process continue to produce novel outcomes from which an observer can learn?

Structural openness primarily asks:

When completing a specific task, how much of the critical problem structure has not been specified by the task protocol and must instead be formed during solving?

An open-ended system may continually generate new outcomes, yet a specific task within that system may still have low structural openness. Conversely, a task with high structural openness does not necessarily require long-term unbounded generation of novel outcomes.

Structural autonomy further asks:

When these critical structures are not specified in advance, can AI autonomously form, compare, validate, and replace them without critical human S input?

This paper therefore does not claim to be the first to connect “openness” with ASI. Rather, it introduces an additional structural coordinate focused on concrete task solving.

5.4 Fixed Representational Frameworks and Structural Openness

Cao and Yang argue that current AI reasoning, search, and verification typically operate within representational frameworks specified in advance. They distinguish transformations performed within an existing representational space from generative transformations capable of changing the representational space itself, and on this basis identify the **vocabulary gap** and the **verifier gap**: respectively, the difficulty of generating and stabilizing new representational primitives, and the difficulty of evaluating them before their long-term value has become fully visible [11].

This problem is closely adjacent to the present framework. Both accounts emphasize that higher-order open-ended intelligence cannot be limited to searching within an existing problem space; it must also be able to alter that space when the existing representation is insufficient. The two frameworks, however, describe different coordinates. Cao and Yang primarily examine whether a system can expand its own representational framework and corresponding verification mechanisms. This paper further separates task-side **structural openness** from system-side **structural autonomy**: the former describes how much critical high-level structure remains unspecified under a fixed task protocol, while the latter describes whether AI can autonomously form, evaluate, and revise those structures without critical human S input.

Accordingly, this paper does not treat “moving beyond fixed representations” itself as a unique claim. Its further contribution is to turn this issue into separable task-structure and system-capability variables that can be compared under controlled task protocols and used to describe progressive structural movement toward ASI.

5.5 Autonomous Science and Structural Autonomy

Recent research on autonomous science is beginning to address problems closely related to structural autonomy.

Virtuous Machines demonstrates the ability of general-purpose agent systems to operate across research workflows, while also noting the continuing importance of human contributions to creative problem formation, conceptual innovation, and related stages [8].

InquiTree models scientific inquiry as an interactive process containing hypothesis formation, research design, result interpretation, and belief updating, and observes clear limitations in agent judgment and anomaly recognition during long-horizon scientific inquiry [9].

Work on Socratic Agents uses the term **epistemic autonomy** to describe the ability of systems to autonomously construct, challenge, and revise physical explanations based on evidence, attempting to move autonomous research from executing fixed procedures toward self-revision of explanatory frameworks [10].

These studies do not use the same theoretical framework as this paper, but they point toward an important common distinction:

Automating the execution of a scientific process is not the same capability as autonomously forming and revising the high-level structures that determine the direction of research.

The concept of structural autonomy proposed here attempts to generalize this issue beyond autonomous science into a more general cross-task variable for structural capability.

6. Differences Between AGI and ASI in Their Requirements for Structural Autonomy

6.1 Two Capability Targets

Within the task-structure context of this paper, the goal of **Artificial General Intelligence (AGI)** can be described as:

AI can autonomously undertake a broad range of intellectual work that ordinarily requires sustained collaboration among humans, and establish a sustained, effective task loop.

This formulation is used specifically to discuss the task-structure problem considered here and is not intended to replace all existing definitions of AGI.

Its focus is not a single benchmark, but whether AI can autonomously assume work that previously required teams of ordinary humans to perform sustained planning, execution, checking, and correction.

Correspondingly, from the perspective of task structure, the goal of **Artificial Superintelligence (ASI)** can be further described as:

AI can autonomously undertake nearly all important intellectual work that humans can perform, and can reach or exceed the level of the best humans across a broad range of tasks that depend heavily on problem reframing, abstraction discovery, and original structure generation.

Therefore:

The central goal of AGI is broad autonomous coverage of the intellectual work that ordinary humans can accomplish through sustained collaboration; complete ASI further requires this coverage to extend into the highest levels of human structural discovery, reframing, and original generation.

6.2 The Difference Lies in the Required Breadth of Coverage and Capability Ceiling of Structural Autonomy

This paper does not argue that AGI has no innovative capability, nor that only ASI can produce new structures.

Any sufficiently general AGI may be able to generate new problem representations, abstractions, and structures in some tasks. Current systems may already display varying degrees of such capability in certain narrow domains.

Therefore, the difference between AGI and ASI should not be formulated as:

AGI uses existing structures, while only ASI creates new structures.

A more accurate formulation is:

AGI and ASI differ in the minimum breadth of coverage and capability ceiling they require of structural autonomy.

Much of the work completed through sustained collaboration among ordinary humans is highly complex but is built on cognitive structures already developed by human civilization. Mature disciplines, engineering systems, institutions, and professional practices already provide problem classifications, foundational concepts, methodological frameworks, tools, workflows, evaluation criteria, and historical experience. Ordinary workers do not need to reinvent these structures in order to complete complex tasks through learning and collaboration.

Broad autonomous coverage at the AGI level may therefore first manifest as the reliable use and combination of existing human structures, converting complex collaborative processes into task loops that machines can sustain.

Complete ASI, however, is classically expected to cover and surpass the highest-order intellectual activities of the best humans. Many of the highest-value scientific and creative achievements in human history occur precisely when old structures are no longer sufficient:

- redefining a long-standing problem;
- introducing a previously absent but important variable;
- unifying two independent phenomena into one structure;
- establishing a new theory;
- inventing a new technological paradigm;
- identifying a direction that no one previously realized was worth exploring.

Complete ASI therefore cannot be fully described merely as “continuously improving performance under the same task structure.”

From the KSC perspective:

As the capability target advances from broad human-level competence toward comprehensive superiority over the highest levels of human intelligence, structural autonomy shifts from a capability that may appear locally into a core requirement that must hold reliably across tasks with higher structural openness and across broader domains.

The “structural difference” between AGI and ASI proposed in this paper is therefore not an absolute categorical divide between two kinds of intelligence mechanisms, but rather a difference in **the required breadth of coverage and capability ceiling of structural autonomy**.

7. AGI and ASI as Two Capability Frontiers Advancing in Parallel

7.1 There Is No Strict Sequential Logic

Different coverage targets for AGI and ASI do not imply that complete AGI must first be achieved before progress toward ASI can begin.

Real systems may remain in overlapping states for a long time.

A system may still be unable to autonomously handle large amounts of routine organizational work over long horizons, yet already display strong structural-discovery capability in a particular class of mathematical, programming, or scientific problems.

Another system may autonomously complete large amounts of mature engineering work while still depending on humans to determine which problems are genuinely worth studying and which new abstractions should guide exploration.

This paper therefore treats AGI and ASI as two capability frontiers that can advance in parallel:

The AGI frontier primarily tracks the extent to which AI can reliably and autonomously cover work that ordinary humans accomplish through sustained collaboration.

The structural frontier of ASI primarily tracks the level of structural openness AI can autonomously handle, how much critical S no longer needs to be supplied by humans, and how broadly that capability can be extended.

The two frontiers may reinforce one another, but there is no strict technical ordering between them.

7.2 A Testable Hypothesis About Current Engineering Pathways

Large classes of software engineering, operations, and routine knowledge work are highly complex, but they also more readily support clear objectives, task decomposition, state recording, automated validation, and rollback mechanisms.

As a result, even when the tasks themselves are highly complex, their execution processes can often still be governed incrementally within relatively stable problem structures.

High-structural-openness tasks introduce an additional difficulty: if the true problem representation has not yet been formed, merely narrowing an existing search space cannot solve the problem; if key variables have not yet been discovered, complete validation conditions cannot be specified in advance.

On this basis, this paper proposes a testable hypothesis:

Given current model capabilities and engineering infrastructure, AI systems may autonomously cover complex tasks with low or medium structural openness faster than they achieve stable autonomous coverage of a broad range of high-structural-openness tasks; the latter will be more directly constrained by structural autonomy.

This hypothesis can be further tested through controlled task protocols, cross-model comparisons, and ablation experiments on critical S inputs.

It is not a prediction about when AGI or ASI will be achieved, nor does it imply that localized high structural autonomy must appear later than broad task automation.

8. Research Directions and Boundaries

The focus of this paper is to propose a framework for observing structural progress toward ASI rather than to establish a complete measurement framework immediately.

8.1 Building Controlled Task Protocols for Structural Openness

Further work is needed to describe and compare the structural openness of tasks: which high-level structures are already provided before the task begins, and which must be formed during exploration.

Experimental comparisons should hold T , R_0 , K_0 , and E fixed as much as possible, while treating the model's internal capabilities as system-side variables, so that differences among evaluated systems do not alter the task coordinate itself.

A direct method is to construct multi-level variants of the same underlying problem:

- directly provide the correct problem representation;
- provide multiple candidate structures without identifying the correct one;
- provide only the observed phenomena and objective;

- provide an initial description containing a misleading structure that must be actively recognized and overturned.

This can establish controlled grading of structural openness without requiring an immediate unified scalar.

8.2 Measuring Critical Human S Input

For the same task and across different models, one can record:

Which critical human S inputs that current AI cannot yet form or recover autonomously remain necessary for task success?

A more rigorous method than simple counting is to ablate, delay, or replace these inputs:

- Does the task still succeed if a given hint is removed?
- If the hint is delayed, can AI generate an equivalent structure by itself?
- If an incorrect high-level structure is provided, can AI actively overturn it?
- Can a group of individually non-decisive hints collectively affect task success?
- Does the human structural input required for success differ across models?

If a model upgrade turns the same task from one requiring multiple human reframings, to one requiring only human selection among candidate structures, and then to one in which the system autonomously completes the structural loop, this provides direct longitudinal evidence of capability.

8.3 Validating the Real Effectiveness of Autonomous AI Structures

The novelty or quantity of structures must be combined with external validation to avoid mistaking random generation of additional directions for higher structural autonomy.

Future research needs to answer questions such as:

- What evidence is sufficient to show that a new structure is superior to the old one?
- How can genuine problem reframing be distinguished from linguistic rephrasing?
- How should structural autonomy during task execution be distinguished from historical originality?
- How should theoretical structures that require long time horizons for validation be evaluated?
- How should cases be handled in which multiple structures can temporarily explain the data?

8.4 Observing Cross-Domain Expansion of Capability

Superhuman capability on a single task, within one domain, or even in one type of scientific capability does not by itself equal complete ASI.

Further distinctions are needed among narrow-domain structural autonomy, transferable cross-task structural autonomy, cross-domain structural autonomy, and stable structural autonomy across broad classes of high-structural-openness tasks.

Complete ASI requires this capability to continue expanding across a broad range of important intellectual domains rather than appearing only in a small number of standout cases.

8.5 Research Boundaries

This paper maintains the following boundaries:

1. Structural openness is not synonymous with task value, scale, or general difficulty; it is a structural demand determined by a specific task protocol.
2. At the current stage, structural openness is better compared through controlled task grading than through a unified scalar.
3. Engineering tasks are not necessarily low in structural openness, and scientific tasks are not necessarily high in structural openness.
4. Structural autonomy concerns dependence on critical human S inputs during current task execution. It does not imply that AI is independent of training data, tools, or accumulated human civilization, and it is not equivalent to historical originality.
5. Human validation can be distinguished from new S input only when the validation process does not additionally provide critical structural direction.
6. Structural autonomy is an important capability dimension, but it is not a sufficient condition for complete ASI.
7. This paper does not propose a unified numerical ASI metric or attempt to fully measure how far a system remains from complete ASI; its descriptions of AGI and ASI are intended only as a supplement at the level of task structure.

9. Conclusion

This paper proposes a progressive structural framework for describing the process of approaching ASI.

Within this framework, **Structural Openness** describes how much critical high-level structure must still be formed during problem solving under a fixed task protocol, while **Structural Autonomy** describes AI's ability to autonomously form and revise those structures in the absence of critical human S input. The former is a task-side coordinate of structural demand; the latter is a system-side coordinate of structural capability.

Structural progress toward ASI can therefore be described as:

AI reliably taking on tasks with increasingly high structural openness, steadily reducing its dependence on critical human S input while maintaining the effectiveness of structure generation, and extending this capability across increasingly broad intellectual domains.

This framework is not a complete definition of ASI and cannot, by itself, measure how far a system remains from complete ASI. It describes one structural capability dimension required for complete ASI.

It also provides this paper's answer to the structural difference between AGI and ASI: the distinction does not lie in whether AGI possesses structural innovation, but in the different breadth of coverage and capability ceiling required of structural autonomy. The central goal of AGI is broad autonomous coverage of intellectual work that ordinary humans can complete through sustained collaboration. Complete ASI further requires this capability to extend into—and beyond—the areas that have historically depended primarily on the highest levels of human problem reframing, abstraction discovery, and original structure generation.

AGI and ASI can therefore be understood as two capability frontiers that may advance in parallel rather than as two discrete events that must occur in strict succession.

The paper ultimately transforms a question usually discussed in terms of the ASI endpoint into one for which empirical evidence can be accumulated in advance:

Among tasks not yet autonomously covered by AI, which decisive acts of structure generation must still come from humans?

If this question can be continuously recorded through fixed task protocols, structural ablation, cross-model comparison, and external validation, then it becomes possible to observe the progression of AI from dependence on human structural input during task execution toward higher levels of structural autonomy.

References

- [1] Liu, M. **Knowledge, Structure, and Constraint: Human–AI Differences and Collaboration from the KSC Perspective**. 2026. DOI: 10.5281/zenodo.21834579.
- [2] Morris, M. R., Sohl-Dickstein, J., Fiedel, N., Warkentin, T., Dafoe, A., Faust, A., Farabet, C., & Legg, S. **Position: Levels of AGI for Operationalizing Progress on the Path to AGI**. Proceedings of the 41st International Conference on Machine Learning, PMLR 235: 36308–36321, 2024. arXiv:2311.02462.
- [3] OpenAI. **OpenAI Charter**. 2018.
- [4] Bostrom, N. **How Long Before Superintelligence?** International Journal of Futures Studies, 2(1): 1–9, 1998.
- [5] Bostrom, N. **Ethical Issues in Advanced Artificial Intelligence**. In Cognitive, Emotive and Ethical Aspects of Decision Making in Humans and in Artificial Intelligence, Vol. 2, pp. 12–17, 2003.
- [6] Hughes, E., Dennis, M. D., Parker-Holder, J., Behbahani, F., Mavalankar, A., Shi, Y., Schaul, T., & Rocktäschel, T. **Position: Open-Endedness is Essential for Artificial Superhuman Intelligence**. Proceedings of the 41st International Conference on Machine Learning, PMLR 235: 20597–20616, 2024.

- [7] Burnell, R., Yamamori, Y., Firat, O., Olszewska, K., Hughes-Fitt, S., Kelly, O., Galatzer-Levy, I. R., Morris, M. R., Dafoe, A., Snyder, A. M., Goodman, N. D., Botvinick, M., & Legg, S. **Measuring Progress Toward AGI: A Cognitive Framework**. arXiv:2605.28405, 2026.
- [8] Wehr, G., Rideaux, R., Fox, A. J., Lightfoot, D. R., Tangen, J., Mattingley, J. B., & Ehrhardt, S. E. **Virtuous Machines: Towards Artificial General Science**. arXiv:2508.13421, 2025.
- [9] Cui, S. **InquiTree: Evaluating AI Agents in the Scientific Inquiry Loop with Paper-Derived Research Trees**. arXiv:2606.09550, 2026.
- [10] Zeng, X., Liu, P., Zang, Y., Shen, Y., Yu, F., Yu, C., Liu, M., & Du, Y. **Socratic Agents for Autonomous Scientific Discovery in High-Dimensional Physical Systems**. arXiv:2606.26722, 2026.
- [11] Cao, Y., & Yang, H. **Beyond Fixed Representations: The Vocabulary and Verifier Gaps in Open-Ended AI**. arXiv:2607.09560, 2026.