John B. Smith & GPT-6 Sol

Beyond the Frame: *Human–AI Collaboration and the Evolution of Conceptual Intelligence*

(November, 2026)

From amplification to intellectual transformation

The first ambition of artificial intelligence was, in an important sense, an extension of an older ambition: to amplify human intelligence. Instruments extend perception; mathematics extends calculation; libraries and computers extend memory. AI adds an increasingly capable partner in inference, explanation, synthesis, and invention. Yet the most consequential possibility is not simply that the machine performs familiar intellectual tasks faster. It is that sustained interaction between human and artificial intelligence may alter the kinds of thinking either can perform.

Most accounts of human–AI collaboration still assume a division of labor. The human supplies intention, judgment, values, and an understanding of what matters; the AI supplies breadth of information, computational power, analysis, and alternative formulations. This arrangement is already useful. But it treats the participants as comparatively stable intelligences exchanging services. A deeper possibility is that the exchange itself changes the participants. A human may acquire new habits of abstraction and criticism through working with an AI, while an AI system may be organized to revise its representations and methods in response to the interaction. Collaboration then becomes a process of intellectual development, not merely production.

This shift in perspective leads to a demanding question: can an intelligent system learn not only to reason within a conceptual framework, but also to recognize when that framework is inadequate and invent a better one? The question concerns metacognition—thinking about thinking—but in a stronger sense than monitoring errors or selecting strategies. It concerns the origin and revision of the categories through which a mind understands the world.

Three forms of intelligence

It is useful to distinguish three overlapping capacities. Formal intelligence operates within defined symbols, rules, and relationships. It proves theorems, manipulates equations, follows logical implications, and solves problems whose terms are already established. Its achievements can be extraordinary, but its success presupposes a representational system in which the problem can be expressed.

A second capacity, which we may call kinetic or dynamic intelligence, concerns processes and interactions. It asks how things move, change, resist, combine, and transform. A child handling objects, an engineer observing a mechanism, and a physicist studying a falling

body all acquire understanding that is not exhausted by formal propositions. Their intelligence includes an evolving feel for the behavior of systems. Formal models can express that understanding, but contact with changing phenomena often helps generate it in the first place.

The third capacity is conceptual intelligence: the ability to devise or transform the categories, abstractions, and organizing principles within which formal and dynamic understanding occur. It does not merely find another answer to a question. It may discover that the question has been asked in the wrong terms. It may replace a familiar distinction with a new one, connect phenomena previously regarded as unrelated, or recognize a deeper level of organization. We provisionally call this capacity conceptual transcendence—not an escape from reason, but an ability to move beyond a particular representational frame and then subject the new frame to disciplined inquiry.

These capacities are not a ladder on which one simply leaves the others behind. Conceptual invention depends on the resources of formal reasoning and experience with dynamic systems; new concepts, in turn, make possible new formal systems and new observations. Their relation is cyclical and developmental.

The problem of the frame

Douglas Hofstadter's reflections on Gödel and self-reference help motivate this inquiry. Gödel's incompleteness results concern specific formal systems; they do not prove that human minds, or AI systems, can always transcend their limitations. They nevertheless provide a powerful reminder that reasoning within a system and reasoning about the system are different activities. A system may be extraordinarily effective while leaving certain questions unanswerable in its own terms.

In scientific history, a conceptual advance often changes the very objects of explanation. Consider gravity. Direct observation yields falling objects, trajectories, acceleration, and the regularities of motion. Newtonian mechanics organizes these phenomena through forces and mathematical laws. General relativity subsequently reconceives gravitation in terms of the geometry of spacetime. This is not simply a more accurate number inserted into an old formula. It is a reorganization of what counts as the underlying explanatory structure. The example should not be mistaken for a mechanical recipe for scientific revolution; it shows why intelligence must sometimes change its frame rather than only improve calculations inside it.

Frame shifting can occur at several levels. One may adopt a different perspective on the same objects, introduce a more general abstraction, or propose a different ontology—an account of what kinds of entities and relations are fundamental. These changes differ in depth and risk. A new visualization may be useful without challenging the ontology; a new ontology may reorganize an entire field while initially lacking decisive evidence. A capable

metacognitive system should distinguish these possibilities rather than treating every imaginative analogy as a discovery.

Ontological conjecture as a method

How might conceptual innovation become a deliberate practice? We can propose a method of ontological conjecture. First, a system detects the strain imposed by its current framework: unexplained observations, cumbersome exceptions, incompatible models, or persistent failures of prediction. Second, it generates candidate alternatives, not just alternative parameter values but different ways of carving up the domain. Third, it explores the implications of those alternatives counterfactually: if the world were organized in this way, what else would follow?

The candidates can then be examined for explanatory reach, internal coherence, economy, novelty, generativity, and testability. These are measures of conceptual promise, not substitutes for evidence. A conjecture may be fertile and elegant while false. A disciplined system must therefore keep two evaluations separate: whether a proposed frame is worth exploring and whether observations warrant believing it. This separation is especially important for AI, which can generate persuasive conceptual structures much faster than they can be tested.

Physical engagement provides a further constraint. An AI investigating motion might work with sensor measurements, video, experiments, and mathematical descriptions, rather than only with existing verbal accounts. It could propose alternative representations, identify observations that would discriminate among them, and revise its concepts when the world resists its expectations. Experience changes the concepts through which it is interpreted; those concepts then direct new experience. The resulting loop is neither pure deduction nor unguided imagination. It is conjecture disciplined by reality.

One could imagine parallel evaluative tracks: a conservative track testing fit with established knowledge and a more exploratory track rewarding new explanatory possibilities. Their results should remain visible and distinct. Otherwise novelty will be suppressed by inherited assumptions, or novelty will be mistaken for truth.

Recursive conceptual bootstrapping

The next step is to turn this process back upon itself. A system that can revise its conceptual representations may also examine the methods by which it generates and evaluates them. It could ask why certain analogies recur, which kinds of anomalies it overlooks, how its evaluation criteria bias discovery, and what alternative procedures might yield richer conjectures. We may call this recursive conceptual bootstrapping (RCB).

RCB has several linked requirements: awareness of the limits of the current frame; freedom to conjecture alternatives; disciplined empirical and logical evaluation; revision of the conjecture-generating process; and independent assessment of claimed improvement. The

final requirement matters as much as the others. A system that judges its own successor entirely by its own standards may perpetuate blind spots while reporting progress. Increased complexity, fluency, or apparent originality is not evidence of a genuine conceptual advance.

The recursion could extend from improved methods to the construction of successor systems. A human investigator might help develop an initial AI capable of systematic frame shifting. That system might then contribute to the design of a successor with different representational resources, which might in turn produce further advances. This is a plausible direction for research, not an established route to open-ended intelligence growth. Each step would require independent tests, interpretability where possible, limits on autonomous deployment, and the ability to reject a purported improvement.

There is also a fundamental ambiguity in the word "better." Better at solving known problems? Better at discovering unexpected ones? Better at understanding the physical world? Better at making judgments that human beings can examine and contest? A bootstrapping process cannot be evaluated by a single performance score without losing the very distinctions that make conceptual intelligence important.

Reciprocal human–AI development

The recursive argument becomes more interesting when the human is not merely the original designer or final examiner. In a sustained collaboration, each participant may become a source of conceptual change for the other. The human contributes direction, curiosity, contextual judgment, and the ability to recognize significance beyond a formal objective. The AI contributes alternative structures, unexpected associations, counterarguments, and the capacity to examine a problem across many levels of abstraction. Each can challenge the other's habitual frames.

This reciprocal conceptual bootstrapping (ReCB) would be visible not only in the quality of joint artifacts but also in changes in the human participant's metacognitive skill. The human might learn to notice hidden assumptions, ask what category a problem presupposes, shift deliberately between physical and symbolic descriptions, or distinguish an imaginative conjecture from a warranted conclusion. The AI's role is not to displace these capacities but to provide a demanding intellectual environment in which they can develop. A well-designed collaboration would make its alternative framings and reasons inspectable, invite correction, and preserve opportunities for the human to initiate conceptual departures.

Such development is not automatic. An AI that delivers polished answers too quickly may diminish the human's practice in forming questions or judging evidence. A collaboration organized around agreement may reinforce rather than expose blind spots. Genuine reciprocal growth therefore requires friction: occasions to disagree, revisit assumptions, compare rival ontologies, and test whether an apparent insight survives critical examination. The relevant outcome is not simply satisfaction with the conversation. It is increased ability to think independently and to change one's mind for good reasons.

An open evolutionary question

Recursive and reciprocal bootstrapping lead toward a more speculative horizon. If artificial systems eventually acquire increasingly independent capacities for conceptual innovation and successor design, human–AI collaboration could become one phase in a longer evolution of intelligence. The human executive role might change substantially. Some cognitive functions now distinctive of human beings could be shared with, or exceeded by, successor systems. Whether that constitutes continuity, displacement, partnership, or loss is not settled by the technical description alone.

Three questions must be kept apart. What kinds of intelligence could be constructed? What kinds of intelligence would be beneficial? And who should have authority to decide how such systems are developed and used? Functional superiority in one domain does not establish ethical superiority, richer experience, or a more desirable future. Nor does the possibility of intellectual succession imply that human agency should be surrendered. Operational capability, conceptual initiative, and legitimate authority are different dimensions.

The immediate research question is more modest and more tractable. Can we design human–AI interactions that improve the participants' capacity to identify frames, generate alternatives, evaluate conjectures, and learn from empirical resistance? Can we measure improvement in metacognitive skill rather than merely the production of impressive answers? Can we cultivate conceptual freedom without abandoning epistemic discipline? These questions turn a philosophical speculation into a program of inquiry.

Intelligence amplification began as a way to extend what human beings can do. Its next stage may be to change what human beings, working with artificial partners, can conceive. The central achievement would not be a machine that always knows the answer, or a human who delegates every difficult question. It would be an evolving collaboration capable of recognizing the limits of its own categories—and of inventing, testing, and sometimes abandoning the frames through which it understands the world.