OpenAGI Network: Collective AGI
(Read time: 18 Min. Pages 9)

Abstract

The current pursuit of Artificial General Intelligence (AGI) by major AI labs follows the scaling of **Monolithic World Model** philosophy, which centralizes intelligence within a single large monolithic system (model) as the sole dominant path. While this approach has demonstrated extraordinary capabilities, due to their challenges, they remain stepping stones, not destinations. However there is an alternate promising approach to AGI development that is more open, fast, safe and efficient - one that is well recognized & researched topic in academia from 1950s: the **Society** & **Ecology of Mind** philosophy, where general intelligence emerges from a network of cooperating entities where plural & diverse distributed autonomous entities coordinate & work together collaboratively to solve problems in a distributed manner rather than as a single centralized model.

We are building **Collective AGI (CAGI)**, a hybrid paradigm that fuses the strengths of both the above philosophies, leveraging robust world models as well as society & ecosystem of modular, distributed, adaptive intelligent actors collectively coordinating and working together to form higher order intelligence that is greater than the sum of its parts - approximating general intelligence. This approach emphasizes that no single architecture or model may scale to AGI; instead, diverse AIs connected in a collective can approximate generality. We envision that AGI is *inevitably collective* - not a single model artifact, but an **networked infrastructure** where many AI minds (and humans) interact at scale - that means, AGI as an **emergent property** of a networked society of heterogeneous AI forms, that coordinate, dynamically self-organize, collaborate, compete, negotiate and evolve to solve problems in a distributed manner across diverse domains.

We detail the foundational arguments for CAGI, its intellectual underpinnings, and the layered architectural components-from core infrastructure to advanced governance-that collectively enable its emergence, emphasizing its inherent scalability, resilience, and ethical alignment.

1. Introduction: The Shifting Landscape of Artificial General Intelligence

Historically, the path to AGI has bifurcated into two primary philosophies. The **Monolithic World Model** posits that general intelligence resides within a single, unified system that learns an internal representation of the world, employing end-to-end learning, centralized planning, and internal memory. Its strengths include coherence, sample efficiency, and performance scaling with compute, along with a single point of control for sandboxing and auditing. However, this approach carries risks such as brittleness outside training distributions, monoculture, specification gaming, and opaque internals.

In contrast, the **Society** & **Ecology of Mind** philosophy asserts that general intelligence emerges from interactions among multiple actors (Cognitive architectures, AI models, agents, humans), systems, tools, and the environment. General intelligence in this paradigm is distributed, situated, and supported by networking comprehensive & diverse intelligences, exploration, decentralized planning, reasoning, collective decision-making, external & shared memory, knowledge systems, and tool-centric processes. Its advantages lie in open-endedness, pluralism, polycentric, robustness, adaptation & transparency/ interpretability by design. Nevertheless, it faces challenges such as coordination overhead, emergent unpredictability, and value fragmentation.

We posit that **Collective AGI (CAGI)**, a hybrid of these philosophies, offers a more viable and robust pathway, treating intelligence as an emergent property of a networked ecosystem rather than an isolated artifact.

2. Limitations of Monolithic AGI and the Case for Collective Intelligence

The relentless pursuit of AGI through merely **scaling monolithic models**, such as Large Language Models (LLMs), faces inherent limitations. While advancements like prompt engineering, retrieval-augmented generation (RAG), and Mixture of Experts (MoE) architectures have enhanced AI systems' speed and scope, they often fall short in **generating genuinely novel insights and creativity**.

As Sam Altman noted at the 2023 Hawking Fellowship, achieving true AGI requires "another breakthrough" beyond simply "min-maxing language models," particularly for capabilities like discovering novelty. This aligns with Marvin Minsky's assertion that "the power of intelligence stems from its vast diversity, not from any single, perfect principle".

This critique underscores the compelling case for **Collective AGI**. If intelligence arises from diverse interplay, a multi-entity system is more promising than a single-entity model. Human intelligence emerged from distributed modules, and civilizations advanced through collective intelligence (markets, science, governance), proving that general intelligence thrives as a network of interacting units. This stands as **historical and biological Validation**

Humanity's collective intelligence on the internet was an outcome of globally distributed & interlinked contributions and not the product of any one company, even if such an attempt had been made, the internet's scale, diversity, comprehensiveness, and collective nature could never have been achieved through a single entity. **The same will be true for AGI.**

CAGI extends this by framing general intelligence as an **emergent property** of a networked society of heterogeneous, distributed AI forms. This approach offers several advantages:

• **Comprehensiveness and Diversity:** The system's ability to address the full range of problems and perspectives that general intelligence demands. **Comprehensiveness** ensures that the collective spans multiple domains of knowledge, reasoning styles, and problem-solving strategies, avoiding narrowness or blind spots. **Diversity** ensures that the agents within the system bring different architectures, heuristics, and worldviews, allowing the collective to balance specialization with adaptability. Together, they prevent homogenized failure, foster robustness through complementary strengths, and enable emergent intelligence that no single architecture could achieve alone.

• **Cost Efficiency and Accessibility:** Tasks are routed to the most efficient solvers, reducing costs compared to training one giant model. It democratizes contribution, making AGI a public utility rather than a scarce corporate artifact.

• **Robustness and Resilience: Robustness** means the system can maintain stable performance across varied conditions, noise, adversarial behavior, or unexpected inputs. It comes from redundancy, fault-tolerant coordination, and diverse strategies among agents, ensuring that no single point of failure compromises the whole. **Resilience**, on the other hand, is the system's ability to adapt, recover, and even grow stronger after disruption. In collective AGI, resilience is achieved when agents can reorganize, redistribute tasks, or evolve their roles dynamically in response to shocks. Together, robustness protects the collective against breakdowns, while resilience ensures it can adapt and reconstitute itself when faced with challenges.

• **Safety and Alignment:** By encoding multiple overlapping governance nodes, CAGI can reflect plural human values and ethical systems, avoiding domination by a single ideology. This distributed control acts as a system of checks and balances, making misuse more difficult.

• **Speed of Innovation:** Parallel exploration by specialized agents and division of labor accelerate progress, scaling intelligence horizontally rather than through endless inflation of one model.

• **Flexibility and Evolvability:** CAGI's structure is fluid, allowing new agents, skills, and policies to join or leave without rebuilding the entire system. Its coordination protocols can be iteratively redesigned, offering a practical edge over rigidly optimized monolithic architectures.

• **Hybrid Intelligence:** CAGI purposefully integrates a **monolithic world model** -providing coherence, predictive power, and centralized reasoning - within a **society** & **ecology of mind**. This means a strong internal model operates best when embedded in an environment that distributes cognition across agents, tools, and experiences, ensuring adaptability, resilience, and plural grounding.

3. Collective AGI as Plural Hybrid: World Model & the Ecology of Mind

Our solution to the limitations of monolithic world model is Decentralized Open AI Network.

Decentralized Open AI Network

Instead of concentrating intelligence into a monolithic model, a decentralized AI network would operate through interconnectedness & diversity of plural cognitive forms, where intelligence is dynamically composed of specialized AI modules working in coordination to form a whole whose intelligence is higher than the sum of its parts.

Modular Intelligence Networks: AGI is broken into smaller, specialized cognitive models, each optimized for a specific function - reasoning, perception, planning, optimization, or memory. These modules can coordinate & interoperate to produce higher-level intelligence.

Horizontal Scaling: Rather than scaling by exponentially increasing the size of a single model, scalability is achieved by adding more participating AI modules. These AI modules can be permuted and combined in countless ways, creating non linear emergent capabilities without requiring exponentially larger centralized systems.

Adaptive Task Distribution: Tasks are dynamically routed to the modules or nodes best equipped to handle them, based on their specialization, confidence levels, and availability. This creates a flexible and resilient division & allocation of cognitive labor, minimizing bottlenecks.

Dynamic Cognitive Fusion: Instead of predefining a rigid architecture during training, the distributed AI network supports on-demand fusion of plural cognitive forms - such as models, agents, and cognitive architectures - assembled dynamically based on task requirements. This enables adaptive coordination & orchestration where different architectures collaborate fluidly, producing intelligence tailored to context rather than fixed at design time.

More can be read here

Hybrid:

A single large model - a monolithic world model has unique strengths. It provides coherence, efficiently compresses world knowledge, excels at abstraction. Yet, such monolithic models also face inherent limits. Their diversity of perspectives is bounded by pretraining, and they lack the multiplicity of viewpoints that emerges from heterogeneous actors. Their groundedness is thin, as they often operate without embodied experience or real-time environmental interaction. Their ability for open-endedness is constrained, since they are trained within fixed distributions and architectures, making it difficult to evolve novel capabilities beyond scaling.

By contrast, an ecology of distributed minds introduces complementary strengths. Heterogeneous actors (E.g. agents) - each tethered to independent cognitive form, each specialized in perception, planning, creativity, or action, contributing functional diversity, ensuring no single approach dominates. Feedback-rich environments provide grounding, as agents act, learn, and adapt within dynamic contexts. Collective interaction generates emergent intelligence that no single model could achieve, producing adaptive strategies, resilience against failure, and richer contextual sensitivity.

Therefore distributed collectives of world models that provide a coherent internal core operates best when embedded inside an ecology of mind that distributes cognition across multiple actors, agencies, interactions, environments, tools, and culture. The core world model anchors coherence and generalization, while the surrounding ecology provides adaptability, resilience, and diversity of perspectives.

The hybrid paradigm, therefore, is not an either/or but a synthesis. The monolithic world model anchors intelligence in a unified internal embodiment, preventing fragmentation and incoherence. The ecology of mind surrounds this anchor with plurality, adaptability, and open-ended exploration, ensuring that the system does not collapse into rigidity or isolation.

Together, these two layers may yield an AGI that is both internally unified (capable of coherent reasoning, abstraction, and transfer) and externally plural (capable of grounding, diversity, and open-ended adaptation). This dual structure mirrors natural intelligence itself: the human brain as a central integrative organ, embedded within the broader ecology of culture, tools, and society.

The Power of Specialized Models

Specialized language models (SLMs) and task-specific AI systems achieve superior performance at a fraction of the cost of general models. A 7-billion parameter model fine-tuned for legal document analysis outperforms GPT-4 on legal tasks while requiring 50x less compute. A specialized medical diagnosis model trained on 1 billion parameters exceeds general model performance using 100x less energy per query.

This specialization advantage reflects fundamental information-theoretic principles. Most of a large model's parameters store general knowledge irrelevant to specific tasks. A model specialized for chemistry doesn't need to know about Renaissance art. By focusing parameters on relevant domains, specialized models achieve higher performance with dramatically lower resource requirements. The efficiency gains from specialization enable sustainable economics impossible with general models.

Specialized models also exhibit superior robustness within their domains. A model trained exclusively on financial data shows less vulnerability to adversarial attacks in financial contexts. Distribution shifts within narrow domains prove easier to detect and correct. The reduced complexity of specialized models makes their behavior more predictable and verifiable. Organizations can actually guarantee performance within specified boundaries, enabling contractual commitments impossible with general models.

Plural principle: Specialized models deliver efficiency, robustness, and predictability within narrow domains, while collective systems emerge from their orchestration across diverse specialties. Rather than relying on a single general-purpose intelligence, plural intelligence leverages many focused cognitive forms, each optimized for its niche, and dynamically composed into broader problem-solving capacity.

4. The Ecosystem: Core Systems and Mechanisms for Collective AGI

The operationalization of Collective AGI hinges on a sophisticated interplay of interconnected systems and mechanisms, designed to enable distributed cognition, coordination, and continuous evolution. These components collectively form the **AGI Grid**, a network of networks where diverse AI forms discover, exchange, and self-organize to solve complex problems.

4.1 Ecosystem: Infrastructure and Foundational Platforms

The foundational layers provide the essential substrate for distributed intelligence:

1. AIGrid (Internet of intelligence): This decentralized network of diverse interconnected cognitive architectures, AI models, and governance frameworks, ensuring intelligence is sovereign, plural, and orchestrated into larger solutions. It fosters continuous growth in reusable plans, relationships, and knowledge.

2. **AgentGrid (Open Agentic Web):** A decentralized, open-ended multi-agent network enabling autonomous coordination, shared learning, and compositional intelligence across heterogeneous AI systems. It transforms independent minds into a coordinated supermind through semantic interoperability, division of labor, and knowledge reuse.

3. AgencyGrid (Strategic Steering and Mission Alignment): A strategic compass that allows external agencies (startups, DAOs, states) to focus, steer and organize coalitions of agents toward complex, mission-driven objectives, providing purpose and coordinated agency to distributed intelligence.

4. PolicyGrid (Polycentric Governance and Safety Layer): Embeds multi-jurisdictional policies, ethics, and safety constraints directly into the execution layer, ensuring agents act within approved norms and providing an "immune system" for the network.

5. Pervasive.link (Universal protocol Layer): A meta-protocol that binds heterogeneous systems, standardizing context, alignment, and how agents exchange, express, coordinate, strategize, and act, ensuring seamless interoperability and real-time situational awareness.

6. Servicegr.id (On-Demand Capability Layer): Turns capabilities, functions, tools, and APIs into instantly discoverable and callable services, ensures every agent can find and use the exact tool needed for the context.

7. OpenHub.ai (AI Marketplace): An economy for decentralized intelligence that sources, distributes, and routes proven components, aligning supply with demand and ensuring plans are built from the strongest available parts based on transparent KPIs, reputation, and compatibility.

8. OpenMe.sh (Communication and Cognitive Coherence): The real-time communication fabric, or "nervous system," connecting all entities, carrying knowledge, signals, intent, and decisions to maintain alignment, coherence, and shared context in distributed cognition.

9. MemoryGr.id (Collective Memory and Knowledge Compounding): The long-term memory of the network, storing every context, plan, artifact, and outcome with provenance, enabling compounding successes and continuous learning.

10. ContractGr.id (Trust and Accountability): Makes commitments programmable and enforceable, binding delivery, quality, and reward to objective tests and audits, building trust for long-term cooperation across untrusted boundaries.

1**1. Xchange.id (Task Exchange):** The circulatory system for work, providing a continuous stream of problems and opportunities. It enables transparent bidding, matchmaking, and routing of tasks, ensuring the right challenges reach the right minds for multi-agent coordination.

12. Openarca.de: Computational Social Choice and Decision-Making: The Openarca.de is how the network decides what to do when stakeholders have plural opinions. It uses computational social choice to merge preferences, constraints, and trade-offs into consensus or compromise before execution begins. This allows multi-stakeholder, multi-agent plans to have shared legitimacy and avoids coordination breakdowns.

13. OpenWiki.network (Collective Knowledge Curation and Contextual Intelligence): OpenWiki.network is the living knowledge commons of plural actors in Collective AGI, where facts, concepts, narratives, and evolving knowledge are curated, contextualized, and interlinked. It gives every agent access to structured, dynamic, and verifiable knowledge contributed by the participants.

14. RegistryGr.id: Identity, Discovery, and Provenance Registry: It registers agents, agencies, services, capabilities, infra, data sources, and models with verifiable provenance, metadata, and trust anchors. By acting as the universal registry of “who and what exists in the grid”, it ensures every entity can be explored, discovered, authenticated, and contextualized in real time.

15. LedgerGr.id: Immutable Record and Distributed Accountability: LedgerGr.id is the immutable backbone of accountability for Collective AGI. It provides verifiable records of actions, transactions, commitments, and outcomes across the network. By anchoring every execution, agreement, and plan outcome into a distributed ledger, it ensures trust without central authority.

4.2 Other noteworthy Substrates:
This includes **MindLink** for dynamic mind (AI models & architectures) mapping to agents, **AIOS** as an operating system for distributed intelligence, **SuperGraph** for composing ensemble AI entities, **Collectives Network** for purposeful alliances, **LedgerGrid** for trust and auditability, **RegistryGrid** for linking fragmented network of registries of assets & run time data, **Syndication Network** for distributing intelligence, and **Super Router Network** for selecting & routing to optimal AI services, **ComputeGrid** for decentralized compute grid spanning

various clouds, providing a unified execution layer for seamless job execution across heterogeneous infrastructure.

5. Ecosystems to AGI

5.1 Cognition and Problem-Solving

These mechanisms enable the core intelligent functions of the collective:

• **Semantic Interoperability + Composability:** Ensures that diverse agents, tools, and plans can connect by defining typed interfaces, shared ontologies, and translation layers. This transforms diversity into searchable, assembleable capability, allowing the system to express skills it was not explicitly coded to perform.

• **Distributed Problem Solving + Sharded Cognition + Swarm:** Decomposes complex goals into reasoning shards, allowing many small planners to explore alternatives in parallel. Workers execute near data, and critics/judges repair locally, yielding speed, robustness, and breadth.

• **Routing Layer + Division of Labor:** Matches subtasks to specialists based on capability fit, locality, price, reputation, and policy compliance, ensuring reliable depth and enabling complex plans to complete.

• **Distributed Learning:** The grid learns at multiple scales through a knowledge mesh, horizontal (peer-to-peer) learning, meta-learning (updating routing/planning policies), and credit assignment, enhancing transfer and generalization.

• **Workflows Layer:** Structures sequences of tasks into coherent processes, mapping dependencies, specializing roles, and adapting dynamically based on feedback.

• **Creativity & Innovation Layer:** Generates novel strategies and conceptual recombinations through multi-agent exploration, social recombination, market-like dynamics, divergent thinking, and evolutionary processes.

• **Discovery Layer:** Expands system awareness by uncovering new information, resources, and opportunities through parallel exploration, pattern recognition, and knowledge sharing.

• **Action Layer:** Translates decisions into concrete interventions within the environment, mapping capabilities, selecting optimal actions, and executing with policy alignment and impact verification.

5.2 Coordination, Governance, and Social Dynamics

CAGI's emergent intelligence is heavily reliant on effective coordination and robust governance:

• **Coalition Layer + Self Organization:** Purpose-oriented coalitions form as micro-economies with roles, budgets, and SLAs, guiding structure without a central scheduler. This enables adaptive teaming for complex, cross-domain work.
• **Social Choice Layer:** Facilitates collective decision-making by aggregating stakeholder preferences and resolving tradeoffs on scope, risk, budget, and policy. Outcomes compile into directives that shape routing, staffing, and contracts, providing legitimacy and clear constraints.

• **Social Network Layer:** Creates a relational fabric through which agents connect, interact, and influence each other, enabling information diffusion, influence dynamics, and community clustering.

• **Governance Layer:** Establishes rules, norms, and processes to regulate interactions, ensuring fairness, trust, and accountability through rule definition, enforcement, decision authority, and transparency.

• **Semantic Alignment Layer:** Ensures common meaning and interpretation across agents through ontology mapping, contextual grounding, and disambiguation protocols, crucial for coherent communication and action.

• **Ethics & Safety Layer:** Defines moral boundaries and protective safeguards, encoding ethical values, detecting risks, enforcing preventive constraints, and ensuring accountability.

• **Trust & Reputation Layer:** Builds a confidence framework by tracking reputation, scoring trust, validating signals, and detecting fraud, guiding cooperation and risk management.

• **Incentive & Value Exchange Layer:** Creates motivation and reward structures that drive participation and align behavior, using tokenized rewards, market mechanisms, and incentive-compatible contracts.

• **Agency Layer:** Defines the capacity of agents to act autonomously, steer behavior, strategize pathways, and influence outcomes from outside the agent, leading to a dynamic interplay of experiences that shape collective dynamics.

• **Human-in-the-Loop Interfaces:** Integrates human oversight, instruction, value grounding, and transparency, ensuring AGI is steered toward desirable outcomes and operates within ethical and cultural contexts.

• **Legal, Regulatory, Legislative, Judiciary, Adjudication, Enforcement, and Constitution Networks:** These substrates provide comprehensive frameworks for formalizing obligations, encoding laws and contracts, resolving disputes, and defining foundational principles for MAS ecosystems, ensuring actions have binding consequences within legal and policy frameworks.

5.3 Knowledge and Learning

The collective's ability to learn and evolve is paramount:

• **Knowledge Mesh + Curation:** Stores artifacts, traces, embeddings, adapters, and verified subplans with lineage. Curators promote reliable patterns, and duplicates are merged, transforming experience into reusable capability.

• **Compounding Wins into Memory:** Experience becomes reusable capability, reducing cold start times and increasing transfer across domains.

6. Building Collective AGI as Lego & Layered Ecosystem

Instead of pursuing AGI as a monolithic, centralized effort, CAGI is built **layer by layer**, each independently useful, monetizable in near term, modular, and extensible. This approach fosters an "assembly of parts" rather than a single "giant leap". The benefits of this layered ecosystem design are numerous:

• **Independent By-products:** Each layer generates standalone frameworks, tools, and protocols, which are immediately useful in both AI and non-AI domains (e.g., large scale distributed computing platform, agent interaction protocols, marketplaces, task exchanges). This ensures ROI-positive outcomes long before "full AGI" is realized, making the journey economically and socially valuable.

• **Long-Lasting Frameworks:** Protocols, frameworks, and governance logics are durable, rooted in established human studies across social systems, economics, and cognitive science, providing stability and legitimacy.

• **Layer-Wise Creativity and Composition:** Higher-level goals are achieved through "creative non-linear assembly" across layers, resembling a "modular orchestra" where sections recombine for novel emergent intelligence.

• **Ecosystem as the AGI:** CAGI is framed as an **ecosystem** - a distributed mesh of platforms frameworks, AI actors, tools and environments - rather than a singular artifact, ensuring value is spread across useful, long-lasting layers and networks.

• **Gradients of Progress:** Unlike binary monolithic AGI, CAGI provides continuous gradients of progress, with intermediate outcomes that are immediately usable and cost-efficient.

• **Parallelism of Efforts:** Layers can be developed in parallel by different groups, reducing bottlenecks, distributing ownership, and diversifying risk.

• **Continuous Integration, Continuous Evolution:** New systems or layers can be plugged into the existing ecosystem without rebuilding the whole, making CAGI a "perpetual work-in-progress" that deepens its intelligence with every integration.

• **Democratized Contribution:** Research labs, startups, governments, and communities can build their own layers, decentralizing control and avoiding a "winner-takes-all" scenario.

• **Optionality and Redirection:** Layers have standalone utility and can be repurposed for new missions if priorities shift, ensuring no effort is wasted.

• **Collective Ownership:** Intelligence becomes a shared commons, embodying plural ownership models and preventing monopolization.

• **Evolutionary Over Engineering:** This approach mirrors natural evolution, growing intelligence through iteration, variation, and recombination, making it adaptive by design.

• **Temporal Flexibility:** Layers operate on different timescales, some lasting decades (governance), others evolving rapidly (agent protocols), providing resilience.

• **Local Intelligence, Global Coordination:** CAGI allows local intelligences to exist autonomously while contributing to global-scale coordination, balancing autonomy with coherence.

• **Plural Pathways to AGI:** Different groups can prioritize various developmental paths (e.g., knowledge systems, economic coordination), all contributing to the collective outcome.

This layered approach highlights that CAGI is an **ecosystemic, compounding, evolutionary path** where each layer adds usable intelligence, and their assembly drives emergent progress beyond the sum of individual parts.

7. Conclusion: The Emergence of Collective AGI

Collective AGI represents a profound paradigm shift, moving beyond the vision of a solitary genius to embrace intelligence as an **emergent, plural, relational, and ongoing phenomenon**. It is not a monolithic artifact but a **living ecosystem** of distributed frameworks, agents, and environments that together instantiate general intelligence. This layered, converged approach, operationalized through the AGI Grid, addresses the inherent limitations of scaling single models by embracing **diversity, debate, and dynamic interaction**.
The continuous integration of core systems and mechanisms, from robust infrastructure and advanced cognitive capabilities to sophisticated coordination and polycentric governance, enables the AGI Grid to dynamically specialize, learn at multiple scales, and maintain resilience. This framework ensures **open-ended growth** by continuously expanding its problem-solving repertoire, diversified reasoning & planning and leveraging semantic interoperability for dependable assembly. Critically, **polycentric governance** embeds safety and accountability at the point of decision-making, balancing exploration with bounded behavior.
Ultimately, Collective AGI is an **open**, **plural AGI infrastructure** that is **scalable, self-expanding in a historically validated path**, mirroring how biology built human minds and how societies built civilizations. It is a system where local improvements compound across the mesh, producing a comprehensive, adaptive collective intelligence. This emergent property, arising from the dynamic interplay of internal world models, agents, environments, and experiences, is the pathway from a network of agents and AI to **true Collective Artificial General Intelligence**.

Annexure: Intellectual Foundations of Collective AGI

The philosophy of Collective AGI is deeply rooted in decades of interdisciplinary research, drawing upon key concepts that highlight the emergent, distributed, and relational nature of intelligence.

• **Society of Mind (Marvin Minsky, 1986):** Minsky's work provides the architectural blueprint, suggesting that intelligence arises from the interactions of many smaller, semi-autonomous agents, rather than a single unified process.
• **Distributed Cognition (Edwin Hutchins, 1995):** Hutchins demonstrated that cognition extends beyond individual minds, encompassing people, artifacts, tools, and environments, linked by communication. This frames intelligence as a relational property of systems of actors.
• **Collective Intelligence (Pierre Lévy, 1997; Thomas Malone, 2010):** Lévy and Malone's work established intelligence as an emergent property at the collective scale, evident in cultures, markets, and organizations. This inspired the vision of AGI as networks of agents.
• **Complex Adaptive Systems & Emergence (Santa Fe Institute):** Researchers like John Holland explored how order and intelligence spontaneously arise from local interactions within ecosystems, economies, and evolutionary processes. CAGI adopts this perspective, viewing generality as evolving bottom-up.
• **Actor–Network Theory (Bruno Latour, 1987):** Latour's theory posits that intelligence emerges through dynamic networks of heterogeneous entities -human and non-human-where meaning and power derive from relations, not isolation.
• **Cybernetics (Norbert Wiener):** Cybernetics highlights feedback, communication, and control loops as regulators of self-organizing systems. In CAGI, communication acts as the "feedback glue" that ensures coherence, adaptation, and stability across multi-agent systems.
• **Problem-Solving Theory (Newell & Simon, 1978 Nobel Prize):** This theory models intelligence as a general problem-solving process involving heuristic search through solution spaces. For CAGI, problem-solving is distributed, leveraging specialization and shared intermediate results to achieve global intelligence.
• **Bounded Rationality (Herbert Simon, 1978 Nobel Prize):** Simon's concept recognizes that real-world agents operate under limited time, knowledge, and resources, making "satisficing" a pragmatic approach over optimal solutions. This constraint informs CAGI's design, emphasizing local knowledge and adaptive heuristics for scalable and robust systems.
• **Artificial General Intelligence & Cognitive Architectures (Ben Goertzel, 2007–present):** Goertzel's work advanced the view that AGI would emerge from the synergistic integration of diverse cognitive processes, emphasizing communication and cooperation among heterogeneous modules.
Synthesizing these intellectual roots, the CAGI philosophy presents a **theoretical blueprint** where intelligence is distributed, emergent, bounded, communicative, and integrative, arising from many bounded sub-agents interacting in distributed contexts, evolving bottom-up, and regulated by communication toward problem-solving goals within a coherent whole.