The Emergent AGI Programme
A Framework for the Emergence of Multi Species-Scale General Intelligence

---

Part I
The Need for a Collective Intelligence Programme

1. Introduction

Artificial intelligence has progressed rapidly over the past decade. Advances in machine learning, large-scale computation, and data availability have produced systems capable of remarkable achievements in language processing, pattern recognition, reasoning, and decision support. These systems now assist with scientific research, industrial automation, creative work, and complex analytical tasks.

Despite this progress, contemporary AI systems remain fundamentally **isolated cognitive artifacts**. Most systems operate as individual models or tightly controlled pipelines within specific organizational environments. They may be extremely capable within narrow domains, but they do not yet form part of a coherent intelligence ecosystem capable of sustained collective reasoning which we believe is crucial for general intelligence.

This limitation reflects a deeper structural characteristic of current AI development. Much of the field remains centered on **model-centric progress**. Researchers and organizations invest enormous effort in training larger and more capable models, often measuring success in terms of benchmark performance or model size. While such efforts have delivered impressive capabilities, they do not by themselves create the conditions necessary for the emergence of general intelligence.

General intelligence requires more than powerful individual systems. It requires **networks of intelligence capable of cooperating, coordinating, and reasoning together**. Human intelligence itself does not exist solely within individual minds; it emerges through interaction among individuals, institutions, and shared knowledge systems. Science, technology, and civilization advance through collective cognition.

A similar principle may apply to artificial intelligence. Rather than emerging from a single monolithic system, general intelligence may arise from **distributed ecosystems of interacting intelligences**. In such ecosystems, many specialized systems contribute perception, reasoning, planning, memory, and evaluation capabilities to larger cognitive processes.

If this perspective is correct, the path toward advanced intelligence requires a shift in strategy. Instead of focusing exclusively on isolated models, the field must develop the **infrastructure and architectures necessary for intelligence ecosystems to form and evolve**.

This paper proposes that the development of such ecosystems should be approached as a **long-term collective intelligence programme**. Like major scientific and technological programmes of the past—such as nuclear energy development, space exploration, or the construction of the internet—this effort must proceed through clearly defined stages of infrastructure and capability development.

The programme described in this paper outlines a four-stage trajectory toward the emergence of network-scale intelligence systems. Each stage introduces new forms of cooperation among intelligent systems, gradually increasing the degree to which intelligence becomes a property of the network rather than of individual components.

2. From Isolated Models to Intelligence Ecosystems

To understand why such a programme is necessary, it is useful to examine the current structure of artificial intelligence systems.

Modern AI systems are often built as **large monolithic models** trained on vast datasets. These models demonstrate impressive capabilities in language processing, visual recognition, and knowledge synthesis. However, even the most advanced models remain fundamentally limited by their architectural design.

A single model must attempt to perform many different cognitive functions simultaneously. It must interpret inputs, retrieve knowledge, reason about relationships, plan responses, and generate outputs. Attempting to concentrate all of these capabilities within one system creates significant challenges in scalability, reliability, and control.

In practice, real-world AI deployments increasingly rely on **compound architectures**, where multiple specialized systems collaborate to complete tasks. A language model may interpret user intent, a retrieval system may access relevant knowledge, a planning module may determine an appropriate sequence of operations, and verification systems may check the accuracy of outputs.

These compound systems already represent a departure from the traditional monolithic model paradigm. They demonstrate that intelligence can be **composed from multiple interacting cognitive components** rather than contained within a single system.

As AI systems continue to proliferate across organizations, devices, and infrastructures, these compound architectures begin to expand beyond individual applications. Systems developed by different groups may interact through APIs, shared data sources, or coordination frameworks. Over time, this process creates **networks of interacting intelligence systems**.

Within such networks, individual systems function as specialized nodes contributing distinct capabilities to larger reasoning processes. Some nodes may excel at perception tasks such as interpreting images or sensor data. Others may specialize in reasoning, simulation, optimization, or planning. Still others may serve as coordination mechanisms that organize workflows among multiple participants.

These interacting systems form what can be described as **intelligence ecosystems**. Within such ecosystems, cognitive capabilities are distributed across many components that interact dynamically. The ecosystem as a whole becomes capable of solving problems that no individual system could address independently.

This shift—from isolated systems to interconnected ecosystems—represents one of the most significant structural transformations in the development of artificial intelligence.

3. Intelligence as a Network Property

When many intelligence systems interact, a new phenomenon begins to emerge. Intelligence is no longer confined to individual models or agents. Instead, it arises from the **structure and dynamics of the network itself**.

This perspective treats intelligence as a **network property** rather than a property of a single computational artifact.

A useful analogy can be drawn from the history of computing. Early computers were isolated machines performing tasks independently. The development of the internet transformed computing by connecting these machines into a global network. Once connected, computers could exchange information, coordinate processes, and support distributed services that were impossible for isolated systems.

A similar transformation may be occurring with artificial intelligence.

When many AI systems become interconnected, they can exchange intermediate results, share knowledge, and coordinate reasoning processes. A problem can be decomposed into sub-problems and distributed among specialized systems. Results from one system can serve as inputs to another, forming chains of reasoning that propagate through the network.

Over time, such networks may evolve into **graphs of intelligence**. In these graphs, each node represents a cognitive system and each connection represents a relationship of reasoning, communication, or collaboration.

These networks may include many different forms of intelligence:

- machine learning models
- reasoning engines

- simulation systems
- knowledge graphs
- planning modules
- optimization algorithms
- human experts

Within the graph, these systems interact continuously, contributing their specialized capabilities to larger cognitive processes.

This architecture creates the possibility of **distributed cognition**, where reasoning processes extend across many systems simultaneously. Hypotheses may be evaluated by multiple participants, simulations may explore possible futures, and verification mechanisms may ensure the reliability of conclusions.

Under such conditions, intelligence becomes an emergent property of the network rather than the product of any single component.

4. Collective Intelligence

The concept of collective intelligence has long been studied in fields such as sociology, organizational theory, and complex systems science. It refers to the ability of groups of individuals to produce insights and decisions that exceed the capabilities of their individual members.

In human societies, collective intelligence arises through institutions, communication systems, and shared knowledge infrastructures. Scientific communities, for example, operate as distributed reasoning networks where researchers build upon each other's discoveries and collaboratively advance understanding.

Artificial intelligence systems may develop similar collective capabilities when they are integrated within sufficiently sophisticated coordination frameworks.

Collective intelligence in artificial systems can emerge through several mechanisms:

- **division of labor**, where specialized systems handle different components of a larger problem
- **knowledge sharing**, where systems access common information resources
- **task coordination**, where specialized systems collaborate on complex workflows
- **verification and evaluation**, where multiple systems cross-check results
- **adaptive collaboration**, where networks self-organize themselves around emerging problems without central control

When these mechanisms operate effectively, the overall system becomes capable of solving problems that individual participants cannot address independently.

However, collective intelligence does not appear automatically when systems are connected. It requires infrastructure that enables communication, coordination, trust, and shared reasoning contexts. Without such infrastructure, networks of AI systems remain fragmented and inefficient.

Developing this infrastructure is therefore a central challenge in the path toward advanced intelligence ecosystems.

5. The Concept of Cognitive Criticality

Large technological systems often pass through critical developmental thresholds where their capabilities change qualitatively rather than incrementally.

In nuclear energy systems, for example, a reactor becomes **critical** when the chain reaction of nuclear fission becomes self-sustaining. Prior to this point, external inputs are required to maintain the reaction. Once criticality is reached, the system produces sufficient energy internally to sustain its own operation.

A similar concept can be applied to intelligence systems.

In early stages of development, networks of AI systems require significant external coordination. Humans design workflows, integrate components, and direct problem-solving processes. The network operates in a **sub-critical state**, where intelligence remains dependent on external orchestration.

As networks of intelligence systems grow in scale and sophistication, they may begin to organize their own reasoning processes. Systems may dynamically discover collaborators, assemble problem-solving structures, and refine their strategies based on feedback from the network.

At intermediate stages, the ecosystem may begin to show early forms of self-propagating cognition. Some reasoning loops may emerge autonomously, collaborative problem-solving structures may assemble without direct instruction, and knowledge generated in one part of the network may trigger exploration and inquiry in another. These intermediate states represent the gradual movement of the system toward a threshold where distributed cognition becomes increasingly self-organizing.

When these processes become sufficiently mature, the intelligence network may reach a state of **cognitive criticality**.

Cognitive criticality can be defined as the point at which a distributed intelligence system becomes capable of sustaining its own reasoning processes without centralized control. At this stage, the network itself generates and propagates reasoning chains, discovers new knowledge pathways, and reorganizes its internal structures to address emerging challenges.

In many discussions about the long-term trajectory of artificial intelligence, a similar threshold is described using the term **singularity** — a moment when intelligence systems begin improving and expanding their capabilities through self-sustaining processes. Within the framework described in this programme, cognitive criticality represents a structural interpretation of that idea.

This transition would mark a fundamental milestone in the development of artificial intelligence. In other words, the point of cognitive criticality in distributed intelligence ecosystems is what may ultimately be recognized as the arrival of the **singularity**.

6. A Programme Approach to Emergent Intelligence

Achieving cognitive criticality aka singularity will likely require sustained development and coordination efforts. The complexity of the challenge suggests that progress must be approached through a **structured & layered programme rather than single sling shot experiment**.

Large technological transformations in the past have often followed staged development strategies. National nuclear energy programmes, for example, progressed through sequential phases involving infrastructure construction, fuel cycle development, and reactor design. Similarly, the development of the internet involved multiple stages of protocol design, network expansion, and application ecosystems.

A comparable approach may be necessary for the development of collective intelligence systems.

The programme proposed in this paper outlines four sequential stages through which intelligence infrastructure may evolve:

1. **Internet of Intelligence** – establishing connectivity and operational infrastructure for intelligent systems.
2. **Open Intelligence Web** – enabling collaboration, coordination, and economic interaction among intelligent actors.
3. **OpenMind** – enabling integrated cognition across distributed intelligence networks.
4. **Emergent General Intelligence** – the stage at which intelligence networks achieve cognitive criticality / singularity and become self-sustaining reasoning systems.

Each stage increases the degree of collective intelligence present within the system. Early stages focus on connectivity and coordination, while later stages enable integrated reasoning and network-scale cognition.

The remainder of this programme paper explores how these stages unfold and how the emergence of collective intelligence infrastructure may ultimately lead to the development of emergent general intelligence systems.