Part II
The Four-Stage Collective Intelligence Programme

7. Introduction: A Staged Path Toward Emergent Intelligence

Large technological transformations rarely occur through a single breakthrough.

Instead, they emerge through **successive stages of infrastructure development**, where each stage establishes the foundations necessary for the next.

The development of modern computing followed such a trajectory. Early systems began as isolated machines performing specialized calculations. These systems were later interconnected through communication networks, eventually giving rise to the global internet. Once this infrastructure existed, entirely new classes of applications—distributed computing, cloud services, and collaborative knowledge platforms—became possible.

A similar pattern can be observed in other major technological programmes. The development of nuclear energy, space exploration, and telecommunications each required multiple stages of infrastructure maturation before their full capabilities could be realized.

Artificial intelligence may follow a comparable path.

Rather than emerging suddenly from a single technological breakthrough, advanced intelligence systems may arise from **progressive development of intelligence infrastructure**.

This paper proposes a **four-stage programme** for the development of collective intelligence systems. Each stage expands the scale and capability of cooperation among intelligent systems, gradually increasing the degree of collective intelligence that emerges from the network.

The stages can be summarized as follows:

1. **Internet of Intelligence** – the infrastructure layer enabling intelligent systems to connect and operate across distributed environments.
2. **Open Intelligence Web** – the ecosystem layer enabling coordination, exchange, and collaboration among intelligent actors.
3. **OpenMind** – the cognitive integration layer enabling distributed systems to participate in shared reasoning processes.
4. **Emergent General Intelligence** – the stage at which intelligence networks become self-sustaining cognitive systems.

Each stage represents a distinct phase in the maturation of intelligence infrastructure.

The transition between stages is not defined solely by technological progress, but by the **degree of collective intelligence that the system is capable of producing**.

In early stages, intelligent systems operate largely independently.

In later stages, networks of systems begin to behave as integrated cognitive structures.

The final stage corresponds to the emergence of **cognitive criticality**, where the network itself becomes capable of sustaining and expanding its own reasoning processes.

8. Stage One — The Internet of Intelligence

The first stage of the programme focuses on the development of foundational infrastructure for distributed intelligent systems.

At present, most artificial intelligence systems operate within isolated environments. Models are deployed within organizational infrastructures, accessed through proprietary interfaces, and managed through tightly controlled operational pipelines. Although these systems may communicate through application programming interfaces, they rarely operate within shared coordination frameworks that allow them to function as participants in a broader intelligence ecosystem.

The **Internet of Intelligence** represents the stage at which intelligent systems become **first-class network participants**.

In this stage, infrastructure is developed that allows models, tools, and intelligent services to connect across distributed environments. Rather than existing as isolated deployments, intelligence capabilities become **network-accessible resources**.

Key capabilities of the Internet of Intelligence include:

- **distributed execution environments** capable of running intelligent systems across heterogeneous infrastructure
- **plug-and-play composability**, enabling intelligent components to be easily integrated, composed, and remixed to rapidly expand capabilities
- **AI networks** that span workloads distributed across multiple environments, allowing systems to collaborate across infrastructure boundaries
- **permissioned and permissionless participation**, enabling intelligent actors to join, interact, and transact across open or governed networks
- **policy and governance layers** that coordinate rules, permissions, and incentives across distributed intelligent systems

The Internet of Intelligence establishes the **operational foundation for intelligence ecosystems**.

However, at this stage the network does not yet exhibit strong collective intelligence.

Systems may be connected, but they remain largely independent in their reasoning processes.

The collective intelligence produced in this stage can therefore be described as **latent collective intelligence**. The network contains many intelligent components, but the structures required for large-scale coordination have not yet fully emerged.

This stage is analogous to the early internet, when computers were first connected through communication networks. Connectivity alone did not immediately produce the collaborative capabilities that would later define the modern web. Those capabilities required additional layers of infrastructure.

The Internet of Intelligence therefore represents the **first step in a much longer developmental trajectory**.

9. Stage Two — The Open Intelligence Web

Once intelligent systems can connect and operate within shared infrastructure, the next challenge is enabling them to **coordinate and collaborate**.

The second stage of the programme introduces the concept of the **Open Intelligence Web**.

In this stage, networks of intelligent actors begin to form **ecosystems of collaboration and exchange**. Systems are no longer merely connected; they are capable of discovering one another, negotiating interactions, and forming cooperative workflows.

Several new forms of infrastructure emerge during this stage.

First, mechanisms for **capability discovery** allow intelligent systems to locate other systems capable of performing specific tasks. Instead of hardcoding workflows, agents can dynamically identify potential collaborators based on their capabilities.

Second, systems begin to participate in **distributed task exchange networks**. In these networks, agents can announce tasks, submit proposals for completing those tasks, and negotiate the terms under which work will be performed.

Third, **capability markets** begin to emerge. Intelligent services—such as models, simulation tools, or data processing pipelines—can be published as reusable capabilities accessible to other participants in the ecosystem.

Fourth, coordination infrastructure enables the formation of **multi-agent workflows** in which multiple systems collaborate to complete complex tasks.

In addition to operational coordination, this stage introduces mechanisms for **knowledge sharing and collective memory**. Shared knowledge repositories allow agents to accumulate insights, store results, and build upon the work of previous participants.

Governance systems also begin to play an important role. As intelligent systems interact across organizational and geographic boundaries, policies must be established to regulate acceptable behavior, ensure safety, and manage access to resources.

Through these mechanisms, the Open Intelligence Web enables the emergence of **coordinated intelligence**.

In coordinated intelligence systems, multiple participants collaborate on problem-solving tasks, but each participant maintains its own internal reasoning processes. The network orchestrates workflows and exchanges results, but cognition remains largely distributed among individual systems.

This stage corresponds to the emergence of **agent societies**, where intelligent actors cooperate through structured interactions.

The Open Intelligence Web therefore transforms the Internet of Intelligence from a network of connected systems into an **ecosystem of collaborating intelligences**.

10. Stage Three — OpenMind

While coordination among intelligent systems is a powerful capability, it does not yet constitute integrated cognition.

In coordinated systems, agents exchange results but continue to reason independently.

To reach the next stage of collective intelligence, systems must move beyond transactional collaboration and begin participating in **shared reasoning processes**.

The third stage of the programme introduces the **OpenMind**.

The OpenMind represents the stage where distributed intelligence systems begin operating as **integrated cognitive architectures**.

In this stage, networks of intelligent systems develop mechanisms that allow them to share intermediate reasoning states, hypotheses, and contextual information. Rather than simply exchanging final outputs, systems contribute to **shared cognitive workspaces** where reasoning processes unfold collaboratively.

Within such workspaces, multiple cognitive modules may simultaneously participate in solving complex problems.

For example:

- perception systems may interpret incoming data
- reasoning engines may generate hypotheses
- simulation systems may evaluate possible outcomes
- verification modules may assess the reliability of conclusions
- planning systems may propose next actions

These components interact continuously, refining each other's contributions and collectively advancing the reasoning process.

The OpenMind architecture introduces several key capabilities:

Shared Cognitive Workspaces

Distributed environments where participating systems maintain common representations of problems, hypotheses, and reasoning states.

Meta-Cognitive Monitoring

Systems capable of observing ongoing reasoning processes and identifying contradictions, uncertainties, or inefficiencies.

Compound Intelligence Architectures

Integrated systems composed of multiple specialized cognitive components working together within a unified reasoning framework.

Dynamic Mind Assembly

The ability to assemble temporary cognitive structures composed of many intelligent participants, tailored to specific tasks.

Through these mechanisms, the OpenMind enables the emergence of **integrated collective intelligence**.

At this stage, intelligence networks begin to function less like loose federations of independent agents and more like **distributed cognitive organisms**.

Reasoning becomes a collaborative process unfolding across multiple systems simultaneously.

11. Toward Cognitive Criticality

As OpenMind architectures expand in scale and connectivity, they may eventually reach a critical threshold.

In early OpenMind systems, reasoning networks remain dependent on external orchestration. Human designers determine how cognitive modules are assembled, which participants contribute to workflows, and how reasoning processes proceed.

However, as intelligence ecosystems mature, new dynamics begin to appear.

Systems may start to autonomously discover collaborators capable of contributing to a problem. They may assemble reasoning structures dynamically, combining perception systems, reasoning engines, and simulation tools into temporary cognitive assemblies.

These assemblies may generate hypotheses, evaluate evidence, refine strategies, and produce insights without direct human orchestration.

As these processes become increasingly autonomous, intelligence networks approach the threshold of **cognitive criticality**.

Before reaching this threshold fully, intermediate states may emerge where parts of the ecosystem begin to exhibit self-propagating reasoning behavior. Certain cognitive loops may form spontaneously, knowledge produced in one domain may trigger new investigative processes in another, and networks of reasoning modules may assemble dynamically in response to complex problems. These intermediate developments signal that the ecosystem is transitioning from coordinated intelligence toward self-sustaining cognition.

At this point, the network produces enough internal cognitive activity to sustain its own reasoning processes. New insights generate new questions, which propagate through the network and trigger further exploration.

Knowledge accumulates through distributed reasoning cycles, and the system continuously reorganizes itself to address emerging problems.

The transition to cognitive criticality represents a profound milestone in the evolution of artificial intelligence systems.

Just as nuclear criticality marks the moment when a reactor becomes self-sustaining, cognitive criticality marks the moment when an intelligence network becomes capable of **sustaining and expanding its own reasoning processes**.

In broader discussions about the future of artificial intelligence, a similar threshold is often described as the **singularity**— the point at which intelligence systems become capable of accelerating their own development and expanding their reasoning capabilities autonomously. Within the framework of this programme, the concept of cognitive criticality provides a structural and infrastructural interpretation of that transition.

The final stage of the programme explores what such systems may look like and how they may reshape the future of intelligence infrastructure.