

1 UNITED STATES PATENT APPLICATION

2 

---

3 TITLE OF THE INVENTION

4 MULTI-AGENT ARTIFICIAL INTELLIGENCE SYSTEM FOR DISCOVERY, ANALYSIS,  
5 GOVERNANCE, AND PARETO-PRIORITIZATION OF NOVEL, HIGH-IMPACT  
6 QUESTIONS

7 

---

8 FIELD OF THE INVENTION

9 The invention relates to machine-learning architectures and, more particularly, to distributed multi-  
10 agent systems that **pro-actively surface “unknown-unknown” questions**, evaluate them on multi-  
11 objective criteria, and output an **auditable, Pareto-optimal frontier** of inquiries for strategic,  
12 scientific, ethical, or philosophical exploration.

13 

---

14 BACKGROUND OF THE INVENTION

15 Progress in every discipline is bounded not by the answers we possess, but by the **questions we have**  
16 **not yet imagined**. Existing AI tools excel at (i) retrieving answers to *known* queries and (ii)  
17 generating questions only as a by-product of answering tasks. These tools lack:

- 18 1. a specialized engine for **detecting conceptual white-space** across heterogeneous corpora;  
19 2. a **multi-agent adversarial/co-operative loop** capable of scoring each new question on  
20 *novelty, impact, feasibility, ethical risk, and cross-domain leverage*;  
21 3. a **transparent governance layer** that can quarantine bio-security or dual-use hazards *before*  
22 public disclosure; and  
23 4. an **immutable, regulator-ready audit trail** that supports reproducibility and compliance.

24 Consequently, enterprises waste resources exploring redundant or low-impact avenues, while  
25 transformative research questions remain undiscovered.

26

---

27 **SUMMARY OF THE INVENTION**

28 The invention remedies these limitations through a **five-component architecture**:

- 29 1. **Question Discovery Agents (QDAs)** generate candidate inquiries by contrasting predictive  
30 gaps in source corpora with anomaly signals from unsupervised models.
- 31 2. **Question Analysis Agents (QAAs)** compute a *score vector*  $\langle \text{novelty, strategic-impact,}$   
32 *feasibility, ethical-risk, cross-domain-leverage*  $\rangle$  for each candidate.
- 33 3. **Question Governance Agents (QGAs)** enforce policy, resolve scoring disputes via  
34 confidence-weighted voting, and **quarantine questions whose ethical-risk exceeds a**  
35 **programmable threshold**.
- 36 4. A **Question Ledger**—an append-only, cryptographically signed record—maintains full  
37 provenance, including agent rationales.
- 38 5. A **Priority Engine** performs a **multi-objective optimization** over the score vectors,  
39 generating a **Pareto frontier** and publishing a rank-ordered queue tailored to user-defined  
40 weights.

41 Key technical advantages:

- 42 1. **True novelty detection** via adversarial generation–analysis loops plus anomaly metrics.
- 43 2. **Policy-aware quarantine workflow** that no prior art discloses.
- 44 3. **Explainability module** that attaches a chain-of-thought digest to every prioritized question.
- 45 4. **Horizontal scalability** through micro-service deployment of agent instances.

46

---

## BRIEF DESCRIPTION OF THE DRAWING

**FIG. 1** (see page 2 of the specification) shows the high-level data-flow among Data Ingestion (101), QDAs (102), QAAs (104), QGAs (105), Question Ledger & Priority Engine (106–108), and the User/Down-stream AI interface (109). Solid arrows denote primary data flow; dashed arrows show governance feedback.

---

## DETAILED DESCRIPTION OF THE INVENTION

### Priority Engine Algorithm

Upon receiving adjudicated score vectors  $S_i$ , the engine solves:

$$\max_{Q_i \in Ledger} f(S_i, W)$$

#### Subject to:

1.  $f$  implements a Pareto frontier search such that no selected question is dominated on all objectives.
2.  $W$  is a user-supplied weight vector enabling dynamic re-ranking without re-analysis.

---

### Quarantine Workflow

If:

$$ethical\_risk(Q_i) > \theta$$

(where  $\theta$  is a policy-defined threshold), the Question Governance Agent (QGA) moves  $Q_i$  to a **quarantine sub-ledger**.

#### Release is subject to:

1. Applicable human oversight policies
2. A revised risk assessment procedure

## 74    **EXAMPLE USE CASES**

- 75        •    **Synthetic-biology red-team:** The system auto-quarantines a question whose lab-protocol  
76            implications exceed a BSL-3 risk score; human bio-safety officers review and approve partial  
77            disclosure.
- 78        •    **Corporate strategy:** Pareto frontier highlights three cross-domain R&D opportunities that  
79            traditional pipeline tools missed, saving 18 months of exploratory spend.
- 80        •    **National-security foresight:** High-risk cyber-warfare questions are flagged and routed to  
81            cleared analysts under sealed audit keys.

## **Abstract of the Disclosure**

### **Title: Multi Agent Artificial Intelligence System for Discovery, Analysis, Governance, and Pareto Prioritization of Novel, High Impact Questions**

A multi agent artificial intelligence system discovers, evaluates, governs, and Pareto prioritizes previously unasked, high impact questions across diverse knowledge domains. Question Discovery Agents generate candidate inquiries; Question Analysis Agents score each inquiry on novelty, strategic impact, feasibility, ethical risk, and cross domain leverage; Question Governance Agents enforce policy, resolve analytic conflict, and quarantine high risk content. An immutable Question Ledger records all artefacts, while a Priority Engine performs multi objective optimization to output a Pareto frontier accompanied by explainable rationales. The architecture surfaces strategic blind spots, embeds ethical foresight, and supplies an auditable foundation for accelerated innovation.

FIG. 1

