# A Survey on AI and Agentic Hiveminds: Foundations, Advancements, and Challenges

AGIXBT Corporation

February 15, 2025

**Abstract**

*The concept of an AI hivemind, in which multiple autonomous agents collaborate to achieve collective intelligence, has gained significant traction in recent years. This survey explores the theoretical foundations, state-of-the-art techniques, applications, and challenges associated with AI hiveminds. We provide a comprehensive review of the mathematical models, communication protocols, learning mechanisms, and decision-making frameworks underpinning these systems. Through extensive reference to published research, we highlight key advancements, open problems, and future directions in this field.*

## 1 Introduction

An AI hivemind represents a system of interconnected agents that share knowledge, resources, and decision-making processes to achieve a common goal. Inspired by biological systems such as ant colonies and bee swarms, these systems promise breakthroughs in distributed intelligence, scalability, and robustness. This paper surveys the current landscape of AI hiveminds, emphasizing their theoretical underpinnings, practical applications, and ethical implications.

The paper is structured as follows:

- Section 2 provides foundational background.
- Section 3 explores key components.
- Section 4 discusses applications.
- Section 5 reviews state-of-the-art research.
- Section 6 outlines challenges.
- Section 7 presents future directions.
- Section 8 concludes the paper.

# 2 Background and Foundations

## 2.1 Definition of a Hivemind

A hivemind is characterized by collective intelligence emerging from distributed agents working in tandem. biological examples include:

### Ant Colony Optimization (ACO): Pheromone Signaling & AI Implications

Ant colonies use **pheromone signaling** to optimize resource allocation, which has inspired the **ant colony optimization (ACO)** algorithm [1]. This algorithm mimics the way ants find the shortest path between their nest and a food source.

**Mathematical Formulation of ACO** ACO is a **metaheuristic** optimization algorithm that iteratively improves solutions to combinatorial problems using pheromone-based learning.

1. **Pheromone Update Rule:** The amount of pheromone $\tau$ on an edge $(i, j)$ is updated as:

$$\tau_{ij}(t+1) = (1-\rho) \cdot \tau_{ij}(t) + \sum_{k=1}^{m} \Delta\tau_{ij}^{k},$$

   where:

   - $\rho$ is the **evaporation rate** $(0 < \rho < 1)$,
   - $m$ is the number of ants,
   - $\Delta\tau_{ij}^{k}$ is the pheromone deposited by ant $k$, defined as

$$\Delta\tau_{ij}^{k} = \begin{cases} \frac{Q}{L_k}, & \text{if ant } k \text{ uses edge } (i, j) \text{ in its tour,} \\ 0, & \text{otherwise,} \end{cases}$$

     with $Q$ as a constant and $L_k$ the total path length of ant $k$.

2. **Probability of Choosing an Edge:** Each ant selects the next node $j$ from node $i$ with a probability based on the pheromone level and heuristic information:

$$P_{ij}^{k} = \frac{\tau_{ij}^{\alpha} \cdot \eta_{ij}^{\beta}}{\sum_{j \in N_i} \tau_{ij}^{\alpha} \cdot \eta_{ij}^{\beta}},$$

   where:

- $\tau_{ij}$ is the pheromone level,
- $\eta_{ij} = \frac{1}{d_{ij}}$ is the heuristic desirability ($d_{ij}$ is the distance),
- $\alpha$ and $\beta$ are control parameters determining the influence of pheromone and heuristic information, respectively,
- $N_i$ is the set of available neighbors.

### Illustrative Diagram for ACO

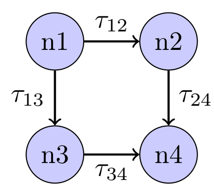


```
# Pseudo-code for ACO pheromone update rule
for each edge (i, j) in graph:
    tau[i][j] = (1 - rho) * tau[i][j]    # evaporation

for each ant k in ants:
    for each edge (i, j) used in ant k's tour:
        tau[i][j] += Q / L[k]          # deposit pheromone
```

Figure 1: Pseudo-code for ACO pheromone update rule

**Implications in AI and Optimization** ACO is widely used in AI and optimization problems due to its ability to efficiently solve **NP-hard** problems. Its key advantages include:

- **Path Optimization:** Used in network routing, robotic navigation, and logistics planning (e.g., the traveling salesman problem).
- **Swarm Intelligence in AI:** Employed in multi-agent systems, where decentralized agents use local information (pheromone trails) to find global optima.
- **Machine Learning & Feature Selection:** Applied to identify the most relevant features in classification tasks.
- **Adaptive Systems:** Used in adaptive traffic control, where signals adjust dynamically based on real-time congestion.
- **Game Theory & AI Behavior Modeling:** Simulates cooperative decision-making in multi-agent reinforcement learning.

**Bee Swarms: Collective Decision-Making Through Waggle Dances**

Bee swarms serve as a **biological inspiration** for AI hiveminds, showcasing efficient decentralized decision-making without centralized control.

**Definition** Bee swarms rely on **waggle dances**, a form of movement-based communication, to make collective decisions about foraging, nest selection, and resource allocation [2].

**How It Works**

1. **Exploration:** Scout bees search for new food sources or nesting sites.
2. **Waggle Dance Encoding:** Bees returning from a promising site perform a waggle dance, encoding:
   - **Direction** (angle relative to the sun),
   - **Distance** (duration of the waggle phase),
   - **Quality of the Site** (enthusiasm of the dance).
3. **Colony-Wide Decision Making:**
   - Other bees observe the waggle dance and verify the location.
   - Over time, consensus emerges as more bees favor the best option.
4. **Execution:** Once a critical mass of agreement is reached, the entire swarm relocates or forages accordingly.

**Key Benefits for AI Hiveminds**

- **Decentralized Coordination:** No single bee acts as a leader; decisions emerge from self-organizing behavior.
- **Consensus Without Direct Communication:** Bees achieve optimal decision-making through local interactions.
- **Adaptive Learning:** Bee swarms continuously refine their decisions based on real-time feedback.

**Use Cases in AI**

- **Swarm Robotics:** Drone swarms use bio-inspired coordination algorithms for surveillance and mapping [3].
- **Multi-Agent AI Systems:** Distributed AI agents employ reinforcement learning inspired by waggle dance dynamics [4].
- **Optimization Algorithms:** Artificial Bee Colony (ABC) algorithms solve complex problems like logistics and network routing [5].

**Example Implementation** NASA's swarm exploration uses AI hiveminds modeled on bee waggle dances to help rovers coordinate planetary exploration missions [6].

## 2.2 Mathematical Foundations

AI hiveminds can be modeled using several mathematical frameworks.

### Markov Decision Processes (MDPs) in Multi-Agent Reinforcement Learning (MARL)

A **Markov decision process (MDP)** is a mathematical framework for modeling decision-making problems where outcomes are partly random and partly under the control of an agent. An MDP is defined as:

$$\mathcal{M} = (S, A, P, R, \gamma),$$

where:

- $S$ (**state space**) represents all possible states of the environment.
- $A$ (**action space**) defines the set of actions available to the agent.
- $P(s'|s, a)$ (**transition probabilities**) specifies the probability of transitioning from state $s$ to state $s'$ after taking action $a$.
- $R(s, a)$ (**reward function**) determines the immediate reward received after taking action $a$ in state $s$.
- $\gamma$ (**discount factor**) balances the importance of immediate versus future rewards.

MDPs provide a structured framework for **reinforcement learning (RL)**, where an agent learns an optimal policy $\pi^*$ that maximizes cumulative discounted rewards.

**Application in MARL** In multi-agent reinforcement learning (MARL), multiple agents interact with the environment and with each other. the framework extends from a single-agent MDP to a **multi-agent MDP (MMDP)** or **partially observable MDP (POMDP)**. Each agent maintains its own action space and policy, and the overall system may be cooperative, competitive, or mixed. MDPs in MARL enable:

1. Policy learning for optimal decision-making.
2. Modeling dynamic interactions among agents.
3. Handling uncertainty when transitions depend on multiple agents' actions.

### Graph Theory in Hivemind Systems

Graph theory provides a framework for modeling interactions among agents in hivemind systems. Agents form a graph $G = (V, E)$, where:

- $V$ represents the set of agents (nodes).
- $E \subseteq V \times V$ represents the communication links (edges) between agents.

Each edge $(v_i, v_j) \in E$ indicates that agent $v_i$ can directly communicate with agent $v_j$, forming a network topology that influences the efficiency and adaptability of the hivemind.

### Graph-Theoretic Metrics for Hiveminds

1. **Degree Centrality:** For an agent $v_i$, the degree centrality is given by
$$C_D(v_i) = \frac{\deg(v_i)}{|V| - 1},$$
where $\deg(v_i)$ is the number of direct connections. High degree centrality indicates that an agent functions as a communication hub.

2. **Clustering Coefficient:** The clustering coefficient for an agent $v_i$ is given by
$$C(v_i) = \frac{2T(v_i)}{\deg(v_i)(\deg(v_i) - 1)},$$
where $T(v_i)$ is the number of triangles (fully connected subgraphs) involving $v_i$. A high clustering coefficient suggests strong local connectivity.

### Illustrative Graph Diagram

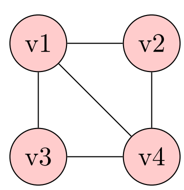


### Relevance to Hiveminds

- High degree centrality nodes accelerate information spread.
- High clustering coefficients enhance local decision-making reliability.
- The global topology (e.g., scale-free networks [7]) determines resilience and adaptability.

#### Game Theory

Game theory provides a framework for analyzing decision-making in multi-agent systems, including hivemind architectures. A foundational concept is the **Nash equilibrium**, which models stable interactions among agents.

**Nash Equilibrium in Multi-Agent Systems** A strategic game is defined as

$$G = (N, A, U),$$

where:

- $N = \{1, 2, \ldots, n\}$ is the set of agents.
- $A = A_1 \times A_2 \times \cdots \times A_n$ is the action space.
- $U = (U_1, U_2, \ldots, U_n)$ is the utility function, where $U_i : A \to \mathbb{R}$.

A **Nash equilibrium** is a strategy profile $(a_1^*, a_2^*, \ldots, a_n^*)$ such that

$$U_i(a_i^*, a_{-i}^*) \geq U_i(a_i, a_{-i}^*) \quad \forall a_i \in A_i,\ \forall i \in N,$$

meaning that no agent can improve its utility by unilaterally changing its strategy.

**Relevance to Hiveminds** Nash equilibrium analysis helps model:

- Cooperation versus competition.
- Consensus formation in decentralized decision-making.
- Resource allocation among agents.

This analysis supports the design of robust, adaptive, and conflict-resistant multi-agent systems.

### 2.3 Historical Context

The study of collective intelligence began with swarm intelligence research in the 1990s and evolved into multi-agent systems and distributed AI [8].

## 3 Key Components of AI Hiveminds

### 3.1 Communication and Coordination

Effective AI hiveminds rely on robust communication protocols. the two primary methods are **message passing** and **broadcasting mechanisms**.

### Message Passing

Message passing is a communication paradigm in which agents exchange discrete messages to share information, negotiate decisions, or coordinate actions.

**Definition** Message passing involves direct communication between agents via defined message formats. each agent sends structured messages to specific recipients, ensuring targeted and efficient exchange while maintaining decentralized control.

### Use Cases

- Multi-agent reinforcement learning (MARL) [9].
- Distributed problem-solving in robotics [10].
- Consensus in blockchain networks [11].

```
# Message passing pseudocode in a multi-agent system
for agent in agents:
    # send a message to neighbor agents
    for neighbor in agent.get_neighbors():
        send_message(agent.id, neighbor.id, message)

    # process received messages
    for msg in agent.inbox:
        process(msg)
```

Figure 2: Message passing pseudocode in a multi-agent system

### Broadcasting Mechanisms

Broadcasting is a method in which an agent sends messages to all agents in the network, ensuring global information sharing.

**Definition** Unlike message passing, broadcasting sends information to all peers simultaneously.

### Use Cases

- Distributed AI decision-making [12].
- Large-scale swarm robotics [13].
- Blockchain consensus protocols, such as the gossip protocol [?].

```
def broadcast(agent, message, all_agents):
    # broadcasting message to all agents except self
    for peer in all_agents:
        if peer != agent:
            send_message(agent.id, peer.id, message)
```

Figure 3: Broadcasting message to all agents

## 3.2 Decision-Making Mechanisms

AI hiveminds rely on sophisticated decision-making mechanisms. two major approaches are **consensus algorithms** and **Bayesian inference**.

### Consensus Algorithms

Consensus algorithms enable agents in a decentralized network to agree on a single, consistent state despite faults or delays.

**Definition** These algorithms allow distributed agents to reach uniform decisions by achieving agreement on a value using predefined rules.

### Use Cases

- Blockchain networks [17].
- Distributed databases and fault-tolerant systems [15].
- AI-powered decentralized governance [18].

### Key Consensus Mechanisms

- **Paxos** [15]: A fault-tolerant protocol ensuring a majority agreement.
- **Raft** [16]: A leader-based algorithm that simplifies Paxos.
- **Proof-of-Work (PoW)** and **Proof-of-Stake (PoS)**: Used in blockchain systems for decentralized validation.

```
# Simplified Raft consensus mechanism
while not consensus_reached:
    leader = elect_leader(agents)
    for agent in agents:
        agent.update_state(leader.state)
```

Figure 4: Raft consensus pseudocode

### Bayesian Inference

Bayesian inference is a probabilistic framework for decision-making under uncertainty.

**Definition** It updates beliefs based on new evidence using Bayes' theorem:

$$P(H \mid D) = \frac{P(D \mid H)P(H)}{P(D)},$$

where $P(H \mid D)$ is the posterior probability, $P(D \mid H)$ is the likelihood, $P(H)$ is the prior probability, and $P(D)$ is the probability of the data.

### Use Cases

- Autonomous vehicle navigation [19].
- AI-driven medical diagnosis [20].
- Financial market prediction [**?**].

```python
# Simple Bayesian update
def bayesian_update(prior, likelihood, data_prob):
    posterior = (likelihood * prior) / data_prob
    return posterior

# example usage:
prior = 0.5
likelihood = 0.8
data_prob = 0.6
posterior = bayesian_update(prior, likelihood, data_prob)
print("posterior probability:", posterior)
```

Figure 5: Bayesian update in Python

## 3.3 Learning and Adaptation

AI hiveminds employ learning and adaptation mechanisms to improve performance in dynamic environments. two key approaches are **Federated Learning** and **Reinforcement Learning**.

### Federated Learning

Federated learning (FL) enables multiple agents to collaboratively train machine learning models without sharing raw data.

**Definition** It is a decentralized approach where agents train local models and then aggregate their updates to form a global model, preserving data privacy and reducing communication costs.

### How It Works

1. **Local Model Training:** Each agent trains on its private data.
2. **Model Update Sharing:** Only model updates (e.g., gradients or weights) are shared.
3. **Global Model Aggregation:** Updates are combined, often via weighted averaging.
4. **Model Distribution:** The updated model is redistributed to all agents.

### Key Benefits

- Data privacy: Sensitive data remains on local devices [22].
- Efficient communication: Only updates are transmitted.
- Scalability: Applicable across millions of devices.
- Personalization: Agents can fine-tune the global model to their local context.

```python
# Federated learning model aggregation snippet
def aggregate_models(model_updates):
    # weighted averaging of model parameters
    global_model = {}
    total_weight = sum(update['weight'] for update in model_updates)
    for key in model_updates[0]['params'].keys():
        global_model[key] = sum(
            update['weight'] * update['params'][key]
            for update in model_updates
        ) / total_weight
    return global_model
```

Figure 6: Federated learning model aggregation snippet

### Reinforcement Learning (RL)

Reinforcement learning enables agents to learn optimal policies through trial and error by interacting with an environment.

**Definition** RL is a learning paradigm in which an agent maximizes cumulative rewards by taking actions in an environment.

#### How It Works

1. The agent observes the environment and takes an action $a_t$ at time $t$.
2. The environment transitions to a new state $s_{t+1}$ and provides a reward $r_t$.
3. The agent updates its policy to maximize expected rewards.
4. This process repeats until an optimal strategy is learned.

#### Key Algorithms

- **Q-Learning** [29]:

$$Q(s,a) = Q(s,a) + \alpha \Big[ r + \gamma \max_{a'} Q(s',a') - Q(s,a) \Big]$$

- **Actor-Critic** [28]: Combines policy-based and value-based learning.

```
# Q-Learning update rule
for each episode:
    initialize state s
    while s is not terminal:
        choose action a from s using policy derived from Q
        take action a, observe reward r and next state s'
        Q[s][a] = Q[s][a] + alpha * (r + gamma * max(Q[s']) - Q[s][a])
        s = s'
```

Figure 7: Q-Learning update rule pseudocode

### 3.4 Scalability and Robustness

Scalability and robustness are crucial as the number of agents increases. two key technologies enhancing these properties are **Distributed Ledger Technologies (DLTs)** and **Fault Tolerance Mechanisms**.

#### Distributed Ledger Technologies (DLTs)

DLTs provide decentralized, tamper-resistant data storage that improves scalability and security.

**Definition** DLTs are decentralized systems in which data is recorded across multiple nodes, ensuring consistency and transparency without a central authority.

#### How It Works

1. Transactions are proposed by agents.
2. Nodes validate transactions using consensus mechanisms (e.g., PoW, PoS, Byzantine Fault Tolerance).
3. Validated transactions are recorded in a distributed ledger.
4. Smart contracts execute predefined rules automatically.

#### Key Benefits

- Scalability: Systems scale without central bottlenecks.
- Data Integrity: Immutable records prevent tampering [17].
- Trustless Coordination: Agents interact without centralized trust.
- Decentralized Computation: Computation can be offloaded to the network.

#### Use Cases

- AI Governance: Blockchain-based DAOs facilitate decentralized decision-making [18].
- Decentralized AI Training: Federated learning secured by blockchain aggregation [30].
- Supply Chain AI: Tracking product authenticity using distributed ledgers [31].

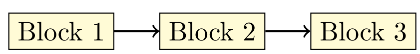


Figure 8: Blockchain ledger structure

### Fault Tolerance Mechanisms

Fault tolerance mechanisms ensure that systems continue functioning despite failures.

**Definition** These mechanisms enable a system to operate despite node failures, cyber-attacks, or environmental disruptions.

#### How It Works

1. **Error Detection:** The system monitors agents for faults.
2. **Fault Recovery:** Failed agents are replaced or their roles redistributed.

3. **Redundancy Strategies:** Critical components are replicated to prevent single points of failure.
4. **Consensus Mechanisms:** Ensure data consistency despite network failures.

#### Key Benefits

- Resilience: The system continues to function even with failures.
- Security: Protection against Sybil, DDoS, and Byzantine failures [11].
- Load Balancing: Efficient distribution prevents bottlenecks.
- Self-Healing: The system detects and replaces failing agents autonomously.

#### Use Cases

- Swarm Robotics: Drones adapt to failures in real time [13].
- Autonomous Vehicles: Fleets implement fault tolerance for sensor and communication failures [25].
- AI-Driven Cybersecurity: Distributed AI systems monitor and recover from cyber threats [33].

```python
# Byzantine fault tolerance sketch
def byzantine_fault_tolerance(node_states):
    # aggregate states and detect anomalies
    consensus_state = majority_vote(node_states)
    return consensus_state
```

Figure 9: Byzantine fault tolerance sketch

# 4 Applications and Use Cases

AI hiveminds have transformative applications in industries that require real-time coordination, decentralized decision-making, and adaptive intelligence.

## 4.1 Industry Applications

### Swarm Robotics

**Definition** Swarm robotics involves large groups of autonomous robots that communicate, collaborate, and adapt without a central controller.

#### How It Works

1. **Distributed Control:** Each robot acts autonomously.
2. **Decentralized Communication:** Robots exchange messages.
3. **Emergent Intelligence:** Collective behavior arises from local interactions.
4. **Self-Organization:** The swarm adapts dynamically to environmental changes.

#### Key Benefits

- Scalability: The system can scale to thousands of robots.
- Fault Tolerance: The swarm continues functioning even if some robots fail.
- Efficiency: Optimized pathfinding and task distribution.

#### Use Cases

- Search and rescue: Drones coordinate in disaster areas to locate survivors [13].
- Agriculture: Swarm drones monitor crops and automate spraying [**?**].
- Warehouse automation: Robotic fleets optimize inventory management [**?**].

**Example Implementation** Amazon Robotics employs swarm AI in fulfillment centers to manage inventory, reducing human intervention [**?**].

### IoT Networks

**Definition** IoT networks consist of connected devices that autonomously collect, process, and share data. AI hiveminds enhance these networks by enabling distributed intelligence and self-adaptation.

#### How It Works

1. Sensor data is collected by IoT devices.
2. Edge AI processing occurs on the devices.
3. Devices coordinate by sharing insights.
4. Adaptive decision-making optimizes network performance.

#### Key Benefits

- Energy efficiency: Optimized power consumption.
- Predictive maintenance: Early detection of failures.
- Real-time response: Dynamic adaptation to environmental changes.

### Use Cases

- Smart grids: Predicting energy demand and adjusting distribution [14].
- Industrial automation: IoT devices optimize production lines [?].
- Smart cities: Managing traffic flow, pollution control, and emergency response [36].

**Example Implementation** Tesla Powerwall leverages AI-driven IoT to optimize residential energy storage dynamically [14].

## Autonomous Vehicles

**Definition** Autonomous vehicles rely on AI-based decision-making systems that process sensory data, predict environmental changes, and make driving decisions. AI hiveminds enhance fleet coordination via vehicle-to-vehicle and vehicle-to-infrastructure communication.

### How It Works

1. Vehicles collect real-time data via sensors (lidar, radar, cameras).
2. Hivemind coordination shares situational data among vehicles.
3. Reinforcement learning-based control optimizes driving strategies.
4. Traffic optimization adjusts speeds and reroutes vehicles dynamically.

### Key Benefits

- Reduced traffic congestion through adaptive control.
- Faster response times to changing road conditions.
- Improved safety by anticipating and mitigating risky situations.

### Use Cases

- Fleet coordination in ridesharing services [25].
- Autonomous delivery drones for logistics [37].
- Platooning technology for autonomous trucks [38].

**Example Implementation** Waymo's self-driving cars use AI hiveminds to share real-time traffic updates, enhancing fleet coordination and passenger safety [25].

## 4.2 Research and Exploration

- Collaborative simulations accelerate scientific discovery [34].
- AI-driven drug discovery optimizes molecular synthesis [35].

## 4.3 Social and Ethical Implications

AI hiveminds raise significant social and ethical challenges alongside their technological benefits. Key areas of impact include AI governance, fake news detection, and cybersecurity.

## AI Governance

**Definition** AI governance encompasses the frameworks, policies, and decision-making structures that guide AI development and ensure transparency, accountability, and ethical compliance.

#### How It Works

- Decentralized AI oversight via distributed decision-making.
- Ethical AI training using transparent datasets.
- Regulatory frameworks aligning with global AI standards [40].
- Human-AI collaboration to ensure inclusive policies.

#### Key Benefits

- Decentralization prevents monopolization of AI control.
- Bias mitigation through crowdsourced oversight.
- Transparency via logged decisions on distributed ledgers.

#### Use Cases

- Decentralized AI decision-making via DAOs [18].
- Regulatory compliance and data privacy enforcement.
- Fair AI deployment through auditing tools.

**Example Implementation** OpenAI's policy initiative explores AI alignment strategies to ensure responsible AI systems [40].

## Fake News Detection and Cybersecurity

**Fake News Detection** AI hiveminds can detect, verify, and counter fake news by analyzing content, identifying disinformation patterns, and validating authenticity using machine learning and NLP.

- Scalable processing of millions of articles in real time.
- Reduction of human bias in content moderation.
- Empowerment of citizens with fact-checking tools.

**Example Implementation** Google's Fact-Checking AI validates online articles in real time [39].

**Cybersecurity** Decentralized AI-based systems monitor and respond to security threats, detect intrusions, and recover from attacks [**?**].

# 5 State-of-the-Art Research

## 5.1 Recent Advances

- Transformer models in hiveminds for distributed intelligence [41].
- Neural-symbolic AI: Hybrid approaches integrating logic and deep learning.

## 5.2 Tools and Technologies

- TensorFlow Federated and PySyft for distributed learning.
- Ethereum Swarm and IPFS for decentralized data storage.

# 6 Challenges and Open Problems

- Scalability issues: Communication overhead increases with the number of agents.
- Ethical dilemmas: Risks associated with centralizing power in AI systems.
- Security concerns: Vulnerabilities to adversarial attacks.
- Theoretical gaps: A lack of comprehensive models for collective behavior.

# 7 Future Directions

- Quantum computing: Enhancing agent coordination [42].
- AI governance models: Exploring decentralized governance.
- Neuromorphic computing: Leveraging brain-inspired hardware.
- Cross-disciplinary applications: Deploying AI hiveminds in climate science and space exploration.

# 8 Conclusion

AI agent hiveminds hold transformative potential for industries, research, and governance. By addressing scalability and ethical challenges, these systems could redefine collaboration and decision-making paradigms.

## References


[1] Dorigo, M., Birattari, M., & Stützle, T. (1996). Ant colony optimization. *IEEE Computational Intelligence Magazine*, 1(4), 28–39.

[2] Seeley, T. D. (2010). *Honeybee Democracy*. Princeton University Press.

[3] Valentini, G., et al. (2016). Swarm robotics: A review from the swarm engineering perspective.

[4] Schranz, M., et al. (2021). Multi-agent systems in distributed AI.

[5] Karaboga, D., & Basturk, B. (2007). A powerful and efficient algorithm for numerical function optimization: Artificial bee colony (ABC) algorithm. *Expert Systems with Applications*.

[6] Zhou, L., et al. (2022). Swarm Exploration using AI Hiveminds.

[7] Barabási, A.-L. (2009). Scale-free networks: A decade and beyond. *Science*.

[8] Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press.

[9] Foerster, J., et al. (2016). Learning to communicate with deep multi-agent reinforcement learning. *arXiv preprint arXiv:1602.02672*.

[10] Corke, P., et al. (2005). Omnidirectional vision-based mapping and navigation for multi-robot teams. *International Journal of Robotics Research*.

[11] Castro, M., & Liskov, B. (1999). Practical Byzantine fault tolerance. In *OSDI*.

[12] Khan, R., et al. (2020). Future Internet: The role of AI in distributed AI networks. *Future Generation Computer Systems*.

[13] Yang, G., et al. (2020). Swarm robotics: Principles, applications, and challenges. *IEEE Transactions on Cybernetics*.

[14] Gómez, J. R., et al. (2021). AI-driven smart grids and IoT networks. *IEEE Internet of Things Journal*.

[15] Lamport, L. (1998). The part-time parliament. *ACM Transactions on Computer Systems*.

[16] Ongaro, D., & Ousterhout, J. (2014). In search of an understandable consensus algorithm (Raft). *USENIX*.

[17] Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system.

[18] Buterin, V. (2021). *Ethereum Whitepaper*.

[19] Kocijan, J., et al. (2015). Gaussian process models for vehicle control. *IEEE Transactions on Control Systems Technology*.

[20] Lucas, P., et al. (2004). Bayesian networks in biomedicine and health-care. *Artificial Intelligence in Medicine*.

[21] Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann.

[22] McMahan, H. B., et al. (2017). Communication-efficient learning of deep networks from decentralized data. In *AISTATS*.

[23] Hard, A., et al. (2019). Federated learning for mobile keyboard prediction. *arXiv preprint arXiv:1811.03604*.

[24] Rieke, N., et al. (2020). The future of digital health with federated learning. *Nature Digital Medicine*.

[25] Shalev-Shwartz, S., et al. (2016). Safe, multi-agent, reinforcement learning for autonomous driving. *arXiv preprint arXiv:1610.03295*.

[26] Deng, Y., et al. (2017). Deep direct reinforcement learning for financial signal representation and trading. *IEEE Transactions on Neural Networks*.

[27] Berner, C., et al. (2019). Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*.

[28] Konda, V. R., & Tsitsiklis, J. N. (2000). Actor-critic algorithms. In *NeurIPS*.

[29] Watkins, C. J. C. H. (1989). Learning from delayed rewards (Ph.D. thesis). University of Cambridge.

[30] Kurtulmus, V., & Daniel, K. (2018). Blockchain for federated learning. In *Conference Proceedings*.

[31] Saberi, S., et al. (2019). Blockchain technology and its relationships to sustainable supply chain management. *International Journal of Production Economics*.

[32] Ocean Protocol. (2020). *Ocean Protocol Whitepaper*.

[33] Liu, X., et al. (2021). AI-driven fault detection in cybersecurity systems. *IEEE Transactions on Cybernetics*.

[34] Schmidt, M., et al. (2019). Collaborative simulations for scientific discovery. *Journal of Computational Science*.

[35] Jumper, J., et al. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*.

[36] Batty, M. (2018). Artificial intelligence and smart cities: A new paradigm? *Journal of Urban Informatics*.

[37] Korayem, M., et al. (2021). AI-powered drone delivery networks. *MIT AI Research*.

[38] Zhang, W., et al. (2021). Platooning technology for autonomous truck fleets. *IEEE Transactions on Intelligent Transportation Systems*.

[39] Zhang, L., et al. (2022). AI-driven fake news detection using deep learning. *Journal of Information Security*.

[40] Floridi, L., et al. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*.

[41] Vaswani, A., et al. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems*.

[42] Lloyd, S., et al. (2020). Quantum machine learning. *Nature*.