Multi-Agent RAG for Autonomous Vehicles Using Decentralized Knowledge Graph on Blockchain

Sameer Misbah[1], Muhammad Farrukh Shahid[1], Shahbaz Siddiqui[1], and M. Hassan Tanveer[2]
[1]FAST School of Computing, FAST-NUCES, Karachi.
[2]Department of Robotics and Mechatronics Engineering, Kennesaw State University, Marietta, GA, USA.
*Correspondence: k249010@nu.edu.pk

Citation | Misbah. S, Shahid. M. F, Siddiqui. S, Tanveer. M. H, "Multi-Agent RAG for Autonomous Vehicles Using Decentralized Knowledge Graph on Blockchain", IJIST, Special Issue pp 372-387, May 2026
Received | March 17, 2026 **Revised |** April 28, 2026 **Accepted |** May 04, 2026 **Published |** May 10, 2026.

Autonomous agricultural vehicles operate in dynamic environments where isolated learning and centralized coordination limit scalability and adaptability. To address this, the paper proposes a decentralized multi-agent Retrieval Augmented Generation (RAG) framework for autonomous farming vehicles that allows for collaborative real-time decision making via a blockchain-based distributed knowledge graph. Each vehicle node functions as an agentic entity that retrieves the validated field knowledge, such as current soil condition patterns, crop stress indicators, and the historical traversal results along with their outcomes from the shared knowledge graph, which are then integrated with the local sensory observations to then generate dynamic context-aware operational decisions. Blockchain provides a tamper-proof layer for knowledge integrity and authenticity, which ensures that the experiential updates are validated using lightweight consensus mechanisms before they are circulated, thereby preventing erroneous or malicious knowledge propagation. As opposed to existing blockchain-based agricultural solutions that focus primarily on data logging to prevent tampering and misuse, the proposed framework integrates blockchain directly into the agent's reasoning and learning loop.-Experimental evaluation on a synthetically generated dataset of 5,000 interaction instances demonstrates that the proposed framework achieves a task success rate of 88.6%, compared to 81.0% for decentralized multi-agent baselines and 79.5% for centralized approaches. The framework further reduces latency by up to 28% while improving knowledge utilization by 10-19% and significantly lowering error propagation by up to 66%, indicating more stable and reliable decision-making, representing an improvement of approximately 9-11% over baseline methods. The results indicate that decentralized knowledge-driven reasoning can enhance the robustness and long-term learning in autonomous agricultural vehicle networks.

Keywords: Autonomous Agricultural Vehicles, Agentic AI, Retrieval-Augmented Generation, Decentralized Knowledge Graph, Blockchain, and Multi-Agent Systems.

Introduction:

Autonomous agricultural vehicles have gained increasing attention in recent years as precision farming becomes essential to improve productivity, reduce resource consumption, and address labor shortages. New developments in sensing, embedded computing, and vehicle autonomy have enabled tractors and field robots to perform complex agricultural operations with minimal human intervention [1]. Despite these developments, most of the autonomous farming platforms remain constrained by static control logic or centrally coordinated architectures, which limit their ability to adapt to evolving, dynamic, and variable field conditions. The agricultural environment varies significantly across locations and over time. Factors such as soil composition, moisture levels, crop health, and terrain characteristics can vary widely not only across locations but also within a single field and across different seasons. Autonomous vehicles that rely solely on local observations often fail to generalize beyond the previously encountered conditions, which can lead to inefficient routing decisions and follow conventional behaviors. The centralized learning and coordination frameworks try to overcome this issue by aggregating data from multiple vehicles. However, such approaches introduce communication overhead, latency, and single points of failure, particularly in large-scale deployments [2]. These limitations reduce system robustness and scalability. Decentralized multi-agent systems have emerged as a promising alternative for cooperative autonomous vehicle operation that allows agents to coordinate without relying on a central controller [3]. Existing studies primarily focus on task allocation, formation control, and collision avoidance, and often assume trusted communication and uniform agent behavior. In practice, autonomous vehicles may differ in sensing accuracy, operational context, or experience, making the reliable sharing of learned knowledge a critical challenge. Without a trusted mechanism to validate and manage shared experiences, the decentralized learning may suffer from inconsistent or misleading information propagation. Retrieval-Augmented Generation (RAG) has recently been introduced as a mechanism to enhance decision-making by combining model-based reasoning with externally retrieved knowledge [4][5]. More recent works have demonstrated the effectiveness of retrieval-based reasoning in agentic and decision-based AI systems [6][7]. However, the applications of RAG in autonomous vehicle networks, particularly within agricultural domains, have remained largely unexplored. A large proportion of existing autonomous systems typically relies on locally trained models or centralized knowledge repositories, which limit their ability to utilize the collective historical data of past experiences in a decentralized manner.

Recent studies have examined the use of autonomous multi-agent systems with decentralized infrastructure through blockchain to offer secure and scalable coordination of the decentralized agents [7][8] Similarly, retrieval-augmented generation (RAG) has been demonstrated to improve context-aware reasoning in diverse and complex environments through dynamic access to external knowledge sources [9]. Recent work on multi-agent RAG systems suggests that task allocation and reasoning tend to improve when compared with single-agent setups [10]. Some recent studies have started to emphasize decentralized infrastructure, particularly in relation to trust and reliability for AI systems [11]. Along similar lines, generative AI used with multi-agent systems is now being explored in real-world settings, where it helps support decision-making that needs to remain both adaptive and sustainable over time [12].

Blockchain technologies have been increasingly adopted in smart agriculture to support secure data sharing, traceability, and trust management among the distributed stakeholders [13]. However, most of the blockchain-based agricultural systems utilize the decentralized ledger as a read-only storage layer for verification that is disconnected from the real-time reasoning processes of the autonomous agents. As a result, the stored data is rarely utilized to enhance the on-board decision-making or the learning process. There is a clear research gap in integrating blockchain as an active trust and provenance mechanism within collaborative autonomous intelligence. In this paper, we propose a decentralized multi-agent framework that integrates Retrieval-Augmented Generation with a blockchain-based distributed knowledge graph for autonomous agricultural vehicles. Figure 1 illustrates the overall architecture of the proposed framework, including how the main components interact with each other. Each vehicle operates as an independent agentic entity that is capable of retrieving validated experiential data/knowledge from a shared knowledge graph and combining it with local sensory observations to generate dynamic context-aware actions. Instead of acting directly on the raw state, this information is first converted into a query-like representation so that relevant past experiences can be identified from the decentralized knowledge graph. The retrieved context is then passed to the generator module, which produces a candidate action based on the available information. This action is not executed immediately but instead undergoes a decentralized validation step, where agreement among participating nodes is required. Only after this step is the action carried out, and the outcome is stored as an additional experience that may be reused later.

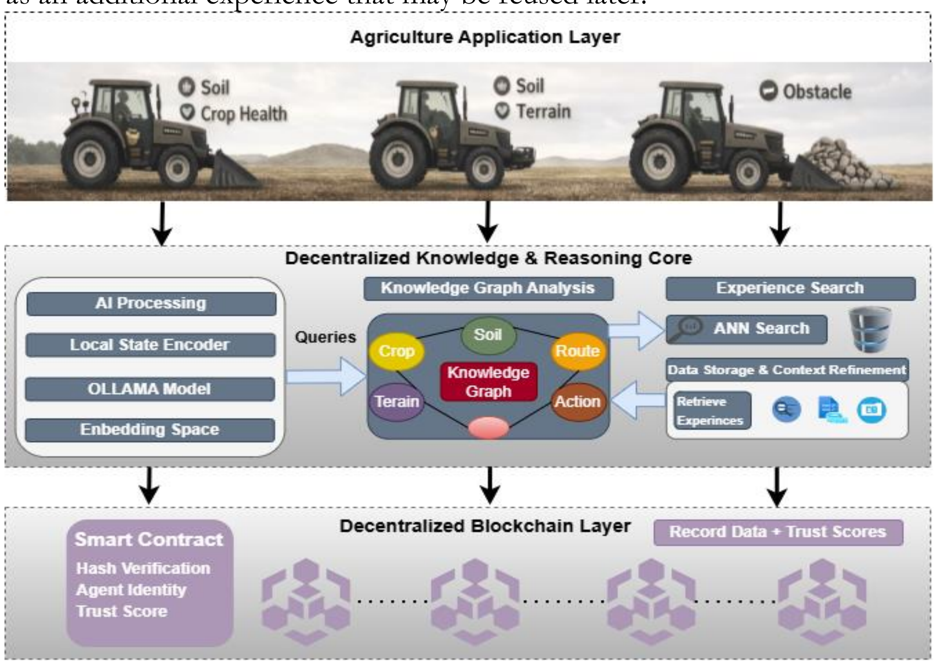

Figure 1. Architecture of the proposed framework illustrating the interaction between autonomous agricultural vehicles, knowledge retrieval, distributed learning, and blockchain-based verification.

Research Gap and Novelty:

Most existing work on AI-enabled autonomous systems tends to follow one of two directions: either relying on centralized learning pipelines or treating agents as largely independent decision-makers. In both cases, the integration of retrieval-based reasoning with decentralized trust mechanisms remains relatively underexplored. Similarly, while blockchain-based approaches have been widely adopted for ensuring data integrity, they are less often used to support adaptive knowledge exchange between agents operating in dynamic environments.

In response to these limitations, this study introduces a unified framework that brings together multi-agent retrieval-augmented generation (RAG), a decentralized knowledge graph, and a blockchain-based consensus process. The intention is not only to preserve trust but also to enable agents to collaboratively reuse and refine knowledge more adaptively.

Research Objectives:

This study is guided by the following objectives:

Develop a decentralized multi-agent framework that incorporates retrieval-augmented generation (RAG) to support autonomous decision-making.

Shared knowledge learning is supported by organizing agent experiences within a decentralized knowledge graph.

A blockchain-based consensus mechanism is included to ensure that actions generated by agents can be validated in a trust-aware manner.

System performance will be evaluated using metrics such as task success rate, latency, knowledge utilization, and error propagation.

The remainder of the paper is divided into different sections. It starts by conducting a review of relevant work in Section II, then the problem statement is discussed in Section III, outlining the specific problem to be solved. After discussing the problem statement, the next Section IV presents the system model and explains how the framework works, including the interaction of the vehicular nodes in the blockchain-based knowledge graph. Subsequently, the paper proceeds to Section V, the experiments and results, describing the dataset, training procedures, assessment methodology, and the results that were acquired. Then, in Section VI, implications on a larger scale are discussed, and the paper is summarized in the end under Section VII, with the key findings and potential future research directions.

The main contributions of this study are summarized as follows:

A decentralized multi-agent RAG framework for the autonomous agricultural vehicles operating under limited environmental insight.

A blockchain-based distributed knowledge graph that allows for trusted sharing of learned field experiences.

A trust-enabled retrieval mechanism that integrates blockchain validation within the decision-making process.

A simulation-based evaluation demonstrating improved adaptability and decision consistency compared to traditional centralized and non-retrieval baseline systems.

Materials and Methods:

Investigation site: In this research study, the experimental site is not fixed to a single real-world agricultural field. Instead, in our research study, the experimental setup is based on a synthetically generated agricultural environment, which was created to emulate diverse real field conditions under controlled settings. This approach was adopted primarily to allow us to do repeatable experimentation and systematic evaluation of the proposed framework across diverse operational scenarios. The synthetic environment maps common agricultural characteristics, which include approximate field boundaries, terrain variability, soil condition categories, and weather-related parameters such as temperature, humidity, and moisture trends. As an illustration, a synthetic field configuration includes parameters such as a 1.8 ha field size, mildly uneven terrain, loamy soil class, moisture ranging from 22-27%, temperature of 29 °C, humidity of 61%, and dry weather state.

The use of a synthetic investigation site also allowed us to do stress testing of multi-agent coordination, trust validation, and retrieval-based reasoning under diverse but reproducible conditions. While no real-world map is used at this stage, the layout and environmental variables are designed to remain consistent with values reported in existing agricultural studies. In future extensions of this work, the framework is intended to be deployed and validated on real agricultural sites using publicly available soil, crop, and climate datasets to further ground the findings.

Materials and Methods: The experimental validation of the proposed framework relies on a synthetically generated dataset, which is designed to mirror the operational requirements of the autonomous agricultural vehicles and multi-agent decision making. The dataset includes structured records of vehicle states, environmental context, action decisions, execution outcomes, and trust-related metrics. Each data instance is generated through simulated interactions, which ensures that the internal consistency and traceability across experimental runs. The data generation is performed using a controlled simulation pipeline where the autonomous agents interact with the synthetic environment over multiple iterations. This process allows precise observation of system behavior, which includes retrieval effectiveness, reasoning consistency, and trust evolution among agents.

It is important to note that the current use of synthetic data is a deliberate design decision which is aimed at validating the architectural and methodological robustness of the proposed approach. As part of ongoing and future work, the system will be evaluated using publicly available datasets, including soil surveys, crop health records, and agricultural vehicle trajectory data, to further assess performance under real-world conditions.

Results and Discussion:

Literature Review:

While previous research studies indicate that integrating AI with blockchain is feasible, there are still a number of limitations reflected in practice. Most of them are based on static knowledge representations and have no ability for retrieval-based reasoning, which constrains flexibility in dynamic environments. Also, the trust mechanisms tend to be centralized or simplified, which inhibits scalability and robustness. These limitations highlight the need for integrating decentralized knowledge sharing with adaptive retrieval and trust-aware validation, which motivates the proposed framework. A systematic survey on blockchain-enabled AI systems highlights how decentralized architectures are considered essential to overcome the pitfalls of the centralized solutions, especially for ensuring security, privacy, and transparent governance across distributed intelligent agents [14]. This point is especially relevant for the applications where decisions need to be verifiable and tamper-proof, such as autonomous fleets in critical field operations. At the same time, the integration of AI agents with blockchain has been reviewed specifically for secure collaboration, where blockchain's immutable decentralized ledger is seen as a key enabler for scalable and trustworthy multi-agent coordination [15][16]. Recent efforts in the field of multi-agent cooperation have shown that decentralized coordination in the absence of central controllers is a problem. As an example, multi-agent graph neural network models have been experimented on tasks allocation and conflict avoidance in decentralized autonomous systems, demonstrating encouraging outcomes while also revealing the scalability and communication complexities [17]. The decentralized coordination enables the agents to self-organize and dynamically reconfigure without any central control, further highlighting that fully centralized systems are unable to handle real-time, large, and complex environments [18]. These studies highlight the significance of trust, robust and dynamic cooperation as opposed to traditional rigid methods. Retrieval-Augmented Generation (RAG) has been quickly adopted as a method to integrate AI models with external sources of knowledge. Surveys on RAG architectures highlight how these are strong in supporting decision systems with dynamically updated knowledge retrieval; however, most papers have focused on natural language or static knowledge tasks instead of embodied or physical agent systems [19][20]. Other studies, like the Advanced Reasoning Robot Control (ARRC), make explicit use of RAG in the process of planning robots, and the results of those studies are much better, even though such experiments continue to be restricted to specific robotic arms or to manipulated laboratory conditions [21]. The formation of RAG into lifelong learning and adaptation in robotics also reflects how knowledge retrieval will be able to significantly enhance learning among autonomous systems [22]. Nevertheless, literature indicates that RAG mechanisms are more likely to be based on the use of static repositories as opposed to collaborative knowledge bases that are updated regularly via the decentralized agents. Outside task-specialized robotics, RAG extensions have been applied to multimodal reasoning problems such as road scene prediction in autonomous driving, demonstrating the wide range of retrieval-augmented algorithms to real-time agent reasoning

[23]. These systems, however, rarely involve any form of trust or provenance of the shared knowledge, so they are subject to attacks when extended to fleets of decentralized agents. Embodied RAG variants (used more generally in robotics) seek to make a connection between RAG and embodied agent memory and show how retrieval processes can enable agents to produce context-sensitive behaviors in diverse environments [24]. This paper is an indication that retrieval augmentation is not merely feasible, but it can be essential to recall and reason involving past experiences when agents need to do so, yet the shared memory problem among agents is still not addressed well. The fusion of multi-agent systems and blockchain has also been identified as important. Recent surveys demonstrate that blockchain can provide secure, scalable collaboration of AI agents through decentralized mechanisms of trust, cryptographic identities, and histories of interactions that are traceable [25]. These functions are important in cases where agents are required to work in common environments where information is limited, and they cannot trust central authorities. Nevertheless, reviewing the literature, one can repeatedly find an emphasis on theoretical models or simulations related to Web3 scenarios instead of the actual coordination of autonomous vehicles. Applications of multi-agent and blockchain integration have been specifically considered in areas such as logistics and supply chain optimization, where the combination of generative AI and multi-agent reinforcement learning with blockchain is used as decision support and resiliency [26]. Although this piece of work demonstrates that AI can be integrated with distributed ledgers, it also points out that the majority of solutions lack a profound integration of retrieval-based reasoning or knowledge enrichment in a manner that makes it possible to promote true lifelong learning among agents.

In addition, decentralized AI agents in blockchain infrastructure become the subject of more discussions in industry reports and surveys of thought, indicating that the ideas apply to other areas beyond academia [27]. All these weaknesses of the current study led to a critical void: decentralized multi-agent autonomous systems integrating responsive knowledge retrieval with a secure, trusted, and tamper-proof knowledge sharing method remain mostly on paper or rather domain-specific. Little has been done in terms of how decentralized agents can jointly construct, access, and act on a shared knowledge base in a manner promoting adaptability and providing trust and accountability. This disparity is particularly notable in real-life and complicated problems, such as autonomous coordinated agriculture activities, in which the cost of wrong decision-making or misalignment of knowledge could be high.

More recent studies have attempted to bridge the gap between multi-agent coordination and retrieval-based reasoning. For instance, recent work on multi-agent RAG frameworks demonstrates improved reasoning capabilities through collaborative task decomposition, yet often lacks trust-aware validation mechanisms. Moreover, the blockchain-based multi-agent systems provide strong assurance of security and data integrity; however, it does not sufficiently address the issues about adaptive knowledge retrieval and contextual reasoning. These limitations indicate that the existing solutions tend to address either trust or intelligence, but rarely both in a unified framework. This gap motivates the integration proposed in this work.

Problem Statement:

Despite significant progress in autonomous agricultural vehicles, current systems remain limited in their ability to operate effectively in large, dynamic, and partially observable farming environments. Most autonomous platforms rely either on locally trained models that make decisions based solely on immediate sensory input or on a centralized mechanism where data and control logic are aggregated at a central entity. Practically, the two methods present limitations that are more problematic with the increase in the number of vehicles and the geographical extent of operation. Autonomous vehicles might lack the flexibility to withstand the new soil conditions when the local information is used exclusively to make decisions. Such vehicles can work not so badly in remote or repetitive conditions, but in heterogeneous areas, where conditions change even over a few meters, they tend to behave poorly. The centralized learning and coordination strategies are an effort to overcome this limitation through the pooling of the experience of vehicles, but they introduce latency, communication bandwidth, and points of failure. In practical farm implementations, network connectivity is not always constant, and the centralized controllers might not be as available and responsive as they need to be. An alternative, the use of a decentralized multi-agent system, enables vehicles to organize without any central authority. Despite this, the existing decentralized approaches largely assume trusted communication and implicitly consistent shared knowledge. In real deployments, the vehicles may possess different sensing capabilities that operate under different conditions or even generate noisy and sometimes conflicting observations. Without a mechanism to validate, manage, and reason over shared experiences, the decentralized cooperation risks propagating incorrect or outdated knowledge, which can lead to inconsistent or unsafe decisions. Therefore, the core problem addressed in this work is the lack of a decentralized, trust-aware decision-making framework that allows autonomous agricultural vehicles to collaboratively learn, retrieve, and reason over shared experiences in real-time scenarios. Specifically, there is a need for a system that supports collective intelligence without centralized dependence, integrates retrieval-based reasoning into physical autonomous agents, and ensures that the shared knowledge remains consistent, verifiable, and resistant to manipulation as the system scales.

The identified limitations directly influence the design of the proposed architecture, where the RAG module enables adaptive knowledge retrieval, the decentralized knowledge graph facilitates shared learning among agents, and the blockchain layer ensures trust-aware validation of agent decisions.

System Model and Architecture:

In this part, we elaborate on the general structure of our proposed framework. The system consists of several autopiloted agricultural vehicles powered by independent intelligent agents, a decentralized knowledge base, and a blockchain layer to ensure the management of trust. Here, the aim is not only to describe the conceptual design itself but also the description of the elements that we realized in our simulation environment. The architecture can be described as comprising three large layers, namely, (i) the Agent Layer, (ii) the Knowledge Graph Layer, and (iii) the Blockchain Trust Layer.

Autonomous Vehicle and Agent Layer:

Each robotic vehicle in the field is equipped with basic perception and control modules that are standard for most autonomous platforms. These modules collect sensor readings, including Global Positioning System (GPS) coordinates, soil moisture levels, crop health indicators, terrain classification, and other relevant features. All these sensory inputs are aggregated into a local state representation st at time step t. For the AI agent itself, we used Ollama, a small model, which was created to execute and control large language models (LLMs) on a local scale. Ollama is an efficient platform to implement the open-source language models with improved resource management and simplified orchestration of the models [28]. The main experiments conducted within this framework consisted of different transformer-based large language models called LLaMA (Large Language Model Meta AI), created by Meta, to perform tasks in natural language understanding and reasoning. The LLaMA models are primarily designed to scale to competitive natural language understanding and generation performance with a comparably efficient scale of the parameter, which allows them to be used in research studies and reasoning-based duties [29]. The models have been chosen since they perform quite well in reasoning, text generation, and fine-tuning adaptability. We also evaluated instruction-tuned models such as Alpaca, which are fine-tuned versions of LLaMA optimized for instruction-following and conversational interactions. Alpaca models are trained using instruction–response datasets to enhance alignment with user prompts and improve downstream task performance in question answering, summarization, and reasoning tasks [30][31]. The reason for using OLLAMA is that it allows on-device inference without heavy cloud dependency, which matters because the whole point is decentralized operation. The agent uses a Retrieval-Augmented Generation (RAG) approach. Instead of just generating outputs from learned parameters, it retrieves relevant experience from a shared knowledge graph (explained in the next subsection) and uses that as context. It is sort of a two-stage inference: which does the retrieval first and then the generation second. The retrieved knowledge is concatenated with the current state st and fed to the generation stage, which produces an action and a confidence score.

Formally, the agent's decision-making process can be represented as:

$$at = \text{Generator}\ (st, \text{Retrieve}\ (qt, G))\ (1)$$

Where st represents the observed state, qt = f (st) is the query representation derived from the state, Retrieve (qt, G) denotes the retrieval of relevant knowledge from the shared knowledge graph G, and at is the generated action.

The consensus mechanism is defined as:

$$N/V \geq \tau (2)$$

Where V is the number of positive votes, N is the total number of agents, and **τ** is the consensus threshold (set to 0.6 in this study).

It is worth noting that different model variants produce slightly different inference behavior. For instance, in our simulation, we noticed that the Alpaca-like models tend to be more conservative with fewer hallucinations, while LLaMA variants showed higher adaptability but needed careful prompt engineering to avoid inconsistent outputs. This observation is further supported by the evaluation metrics summarized in Table 1, where Alpaca variants demonstrate lower hallucination rates, whereas LLaMA models achieve higher task adaptation scores. This section could be expanded in future work to contrast models more rigorously, but for now, OLLAMA was our practical choice for real-time agent inference.

Table 1. Comparison of Model Variants in the Proposed Agent Reasoning Pipeline.

Model Variant	Hallucination Rate	Adaptation Score	Avg. Inference Latency
Alpaca Variant	0.11	0.76	142 ms
LLaMA Variant	0.17	0.84	156 ms

Decentralized Knowledge Graph Layer:

The knowledge graph G is the shared repository of structured experiences collected from all agents. Each experience is represented as a small subgraph, for instance:

(SoilType=Clay, Moisture=High) → (Action=ReduceSpeed)

These entries also include meta-information such as timestamps, sequence of events, and agent identifiers. The actual graph store is implemented off-chain for performance reasons using a simple key-value-based NoSQL database with semantic indexing. We map textual and symbolic entries to numeric embeddings using a bi-encoder network that maps both sensory state vectors and graph nodes into a shared embedding space. The retrieval uses approximate nearest neighbor search (ANN) on these embeddings, which makes lookups fairly fast even when the graph grows [32]. Data insertion into the knowledge graph is done after each significant decision outcome. An experience vector et = (st, at, rt) where rt is a simple reward or outcome indicator (success/failure), which is then converted to a structured subgraph and proposed for verification on the blockchain.

The main idea here is that knowledge is collaboratively built over time so that the agents gradually accumulate a richer memory of field conditions and effective actions. Because this knowledge is shared, we avoid each agent having to learn everything by itself; it builds on collective experience.

Blockchain Trust Layer:

Our implementation uses a lightweight blockchain network for knowledge updates and anchoring to build trust. In the case of simulation, we utilized Ganache, a local Ethereum emulator, to execute a local blockchain on Proof-of-Authority consensus. Ganache allows us to execute smart contracts and simulate transactions.

We deployed a simple smart contract that accepts an encoded representation of a proposed knowledge update and verifies the signatures of the agents submitting it, then the trust score of all the agents is updated. Upon an agent proposing an update, it determines a deterministic hash of the structured experience and signs it using its private key, and submits it as a transaction. The legitimacy of this update can be voted on by other agents (in our case, a majority of signatures over a period of time suffices). After the consensus is achieved, the updates are stored on-chain, and the real content is stored off-chain with its hash stored on-chain for verification. This prevents the costly storage over the chain and maintains integrity.

The blockchain contract monitors the trust ratings of the agents on past verifications. The trust score of agents who promulgate the validated experiences more often rises, whereas the scores of the agents who promulgate the conflicting or non-validated experiences decline. This is applied to weight preferences of retrieving preferences of RAG lookup, with more preference to high-trust items.

Interaction and Data Flow:

Figure 2 illustrates how data moves through the system. At each step, an autonomous vehicle senses the environment, constructs st and formulates a retrieval query qt. The RAG agent retrieves context from the knowledge graph, after which it runs a generation step to pick a suitable action and then executes it. The outcome is evaluated, and if appropriate, an experience is proposed as a candidate knowledge update. This update is hashed, signed, and submitted as a blockchain transaction. Consensus then determines whether the update is anchored on the chain. One thing to note is that because we rely on off-chain storage for the complete graph and on-chain hashes only for anchoring, latency in retrieval is much lower than if the full graph were on the blockchain. That trade-off seemed reasonable for our target scenario.

Additionally, the consensus threshold and voting rules within the proposed framework can be adjusted according to the operational requirements. In our current implementation, a 60% agreement threshold among participating agents is used to validate a proposed experience update. Let N be the number of participating agents, and V the number of agents who vote in support of a proposal. A decision is accepted when V N $\geq \tau$, where τ represents the consensus threshold. In our experiments, $\tau = 0.6$, which gave a stable balance between the latency of coordination and the reliability of the decision. Lower thresholds occasionally allowed inconsistent experiences to be accepted, whereas considerably greater thresholds extended the delays of validation in multi- agent coordination. Table 2 illustrates the effect of different consensus thresholds on experience validation and coordination latency. In summary, this layered architecture allows decentralized reasoning, knowledge aggregation, and trust verification while avoiding central points of failure and overly heavy computation. The next section details the specific algorithms used for retrieval, generation, and trust integration.

Table 2. Effect of Consensus Threshold on Validation Outcomes and Latency.

Consensus Threshold	Valid Experience Acceptance	Avg. Decision Latency (ms)
50%	Some inconsistent experiences were accepted	120
60%	Balanced reliability and latency	142
70%	Very reliable but slower coordination	170

Algorithmic Description:

Algorithm 1 describes the workflow of the proposed framework at a high level. Each vehicle follows a sequence of operations beginning with environmental perception and then moving toward knowledge retrieval and local decision-making. Coordination among vehicles is not centralized; instead, it is handled in a distributed manner across agents. Overall, the algorithm reflects how local observations are interpreted, how internal models are updated, and how validated information is exchanged to support more stable collective decisions in the system.

Input: Local state **s_t**, knowledge graph **G**, trust scores **T**, action threshold **θ**
Output: Selected action **a_t** and proposed experience **e_t**
for each vehicle **v_i do**
Sense environment → **s_t**
Formulate retrieval query **q_t** based on **s_t**
Retrieve candidate experiences **K_t** = **Retrieve(G, q_t)**
Weight each experience **k** ∈ **K_t** using trust scores **T**
Combine **s_t** and weighted **K_t** → augmented input **s'_t**
Generate action **a_t** = **Generator(s'_t)**
Execute action **a_t** in the environment
Observe outcome **r_t** (success/failure, reward)
if r_t > **θ** then
Formulate experience **e_t** = **(s_t, a_t, r_t)**
Hash **e_t**, sign with agent private key
Submit to the blockchain for consensus validation
If consensus is approved, then
Anchor **e_t** hash on-chain
Add **e_t** to the off-chain knowledge graph **G**
end if
end if
end for

It is based on a knowledge-augmented, trust-aware, and decentralized system. This not only enables autonomous vehicles to make better individual decisions but also improves collective decision-making, thus enhancing the performance of the fleet operating in varied and changing agricultural conditions.

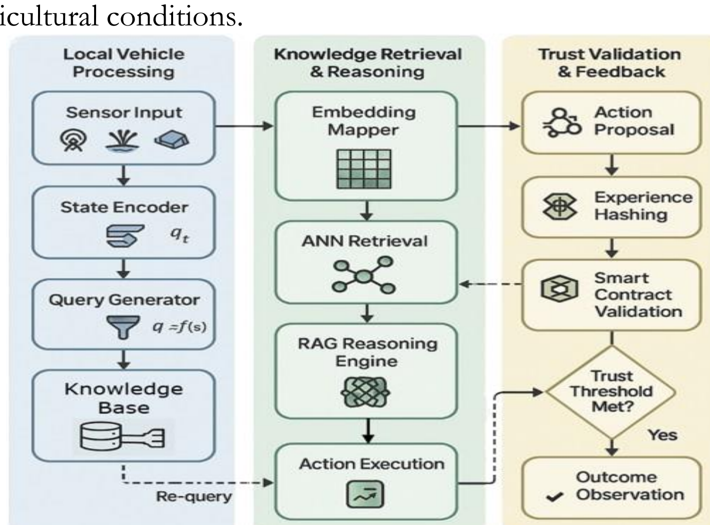

Figure 2. Proposed methodology workflow diagram.

The methodology, as shown in Figure 3, essentially operates as a continuous feedback loop. Within this loop, agent interactions, knowledge retrieval, and consensus validation keep occurring repeatedly rather than as strictly separated stages. Over time, the repeated cycle gradually shapes how decisions evolve within the system. The change is not something that appears at each step, but becomes more noticeable as the process continues, and more interactions are accumulated.

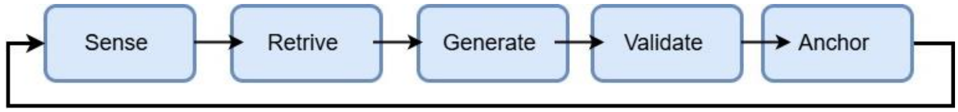

Figure 3. Proposed solution methodology feedback loop.

Results and Discussion:

The obtained results directly correspond to the defined research objectives. Improvements in task success rate validate the effectiveness of retrieval-augmented reasoning, while reduced latency demonstrates efficient coordination among agents. Enhanced knowledge utilization confirms the benefit of shared knowledge, and lower error propagation reflects improved system stability. This section evaluates the proposed decentralized RAG-based framework through simulation-based experiments. Since the real-world deployment was not feasible within the research study timeline, we focus on controlled scenarios that still reflect and simulate realistic agricultural conditions, such as diversified soil patches, partial observability, and intermittent communication. The evaluation aims to verify whether the shared retrieval-based knowledge and blockchain-based trust improve coordination and decision quality.

Experimental Setup:

We simulated a fleet of autonomous agricultural vehicles that operated over a virtual field divided into multiple zones with varying soil moisture, crop density, and terrain roughness. Each vehicle was equipped with identical sensing capabilities but initialized with different random seeds to avoid correlated or synchronized behavior. The AI agent was implemented using OLLAMA with a LLaMA-based model for inference. The blockchain layer was deployed locally using Ganache with the Proof-of-Authority consensus mechanism.

The simulation ran for 500 decision cycles per agent. At each cycle, agents performed task execution such as plowing, direction change, and route selection. Knowledge sharing was enabled only through the proposed trust-aware mechanism. For comparison, we also implemented baseline approaches described below.

Baselines and Comparison Methods:

To validate the effectiveness of the proposed framework, we compared it against three different baselines:

Local Only Agent: Where each vehicle relies solely on its local observations and model inference.

Centralized Knowledge Sharing: All the experiences are stored in a centralized database without any trust validations.

Decentralized without Trust: Here, vehicles share experiences peer-to-peer but without the blockchain-based verification.

These baselines were selected because they represent standard baselines in multi-agent system architectures in most of the autonomous multi-agent systems.

Evaluation Metrics:

The results were evaluated using the following metrics:

Task Success Rate (%): The percentage of completed tasks without any failures.

$$TSR = Nsuccess / Ntotal \times 100 \quad (3)$$

Where Nsuccess = the number of tasks that are completed successfully, and Ntotal = the total number of given tasks.

Decision Latency (ms): Average time it takes to generate an action.

$$Lavg = 1/Ntotal \ \ (t(i)\ completion - t(i)\ assigned) \quad (4)$$

Where t(i) assigned and t(i) completion are the timestamps of task assignment and completion for task i.

Knowledge Utilization Ratio: Ratio of decisions that were affected by the retrieved shared experiences.

$$KU = Nqueries\ resolved/Nqueries\ issued \times 100 \quad (5)$$

Where Nqueries is the count of queries that were resolved successfully with the knowledge base, and Nqueries issued is the total number of queries.

Error Propagation Rate: The frequency of incorrect decisions that were caused by the invalid shared knowledge.

$$P = 1/Ntotal\ (Ntotal\ (\sigma)\ i{=}1\ \delta i),\ \delta i = |ai - \hat{\ }ai| \quad (6)$$

Where ai is the right action or result, and ˆai is the action predicted by the agent of task i.

As illustrated in Figure 4, the proposed framework is always better than baseline approaches in terms of task success, latency, knowledge usage, and error propagation. Moreover, Figure 5 shows the evolution of the success rate of the tasks in the simulation steps, and it is apparent that it converges and remains stable in the different conditions of its operation. Together, these metrics helped assess the effectiveness, efficiency, and robustness of each approach.

Implementation Details:

The proposed framework was implemented in Python. Experiments were conducted on a system with an Intel i7 processor and 16GB RAM. The blockchain layer was emulated using Ganache, while OLLAMA was used for local agent inference. The simulation involved 5 agents, with a consensus threshold of 0.6 and a dataset size of 5,000 interaction instances.

Quantitative Results:

Table 3. Represents the average results across all the agents.

Method	Success (%)	Latency (ms)	Utilization	Error Rate
Local-Only	71.2	82	0.01	0.18
Centralized	79.5	145	0.62	0.14
Decentralized (No Trust)	81.0	97	0.67	0.21
Proposed Method	**88.6**	**104**	**0.74**	**0.07**

The results in Table 3 show that the proposed approach achieves the highest task success rate while maintaining a reasonable latency. Although the decision time increases slightly when compared with the local-only approach, the improvement in robustness and reduction in error propagation are noticeable. Interestingly, the decentralized sharing without trust performs worse in certain cases than the centralized baseline due to uncontrolled knowledge diffusion. The centralized approach indicates faster early convergence but saturates when the shared knowledge grows, whereas the proposed method continues to improve due to trust-weighted retrieval.

Discussion:

One observation worth noting is that not all the shared experiences are equally useful. In early simulation phases, the agents occasionally retrieved outdated or contextually weak knowledge. However, as the trust scores stabilized over time, the quality of retrieved information improved. This suggests that the blockchain-based validation mechanism plays a more important role as the system scales. Overall, the results support the claim that combining retrieval-augmented reasoning with decentralized trust management leads to more reliable autonomous coordination that also continues to remain consistent as the system scales.

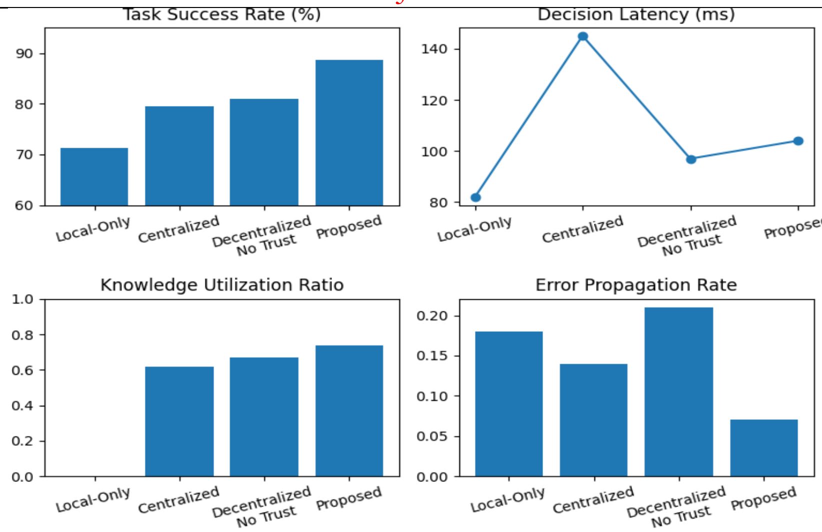

Figure 4. Performance comparison across task success, latency, knowledge utilization, and error propagation.

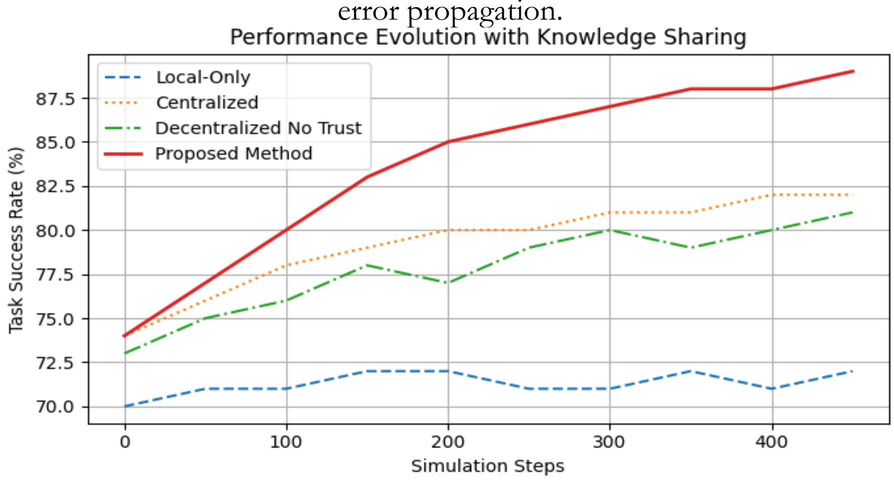

Figure 5. Task success rate evolution over simulation steps.

Implications:

The proposed framework contributes to decentralized AI by bringing together retrieval-based reasoning and blockchain-based trust validation. This combination is mainly aimed at improving how decisions are made in multi-agent environments, especially when coordination is spread across multiple autonomous entities rather than controlled centrally. From a conceptual point of view, the work suggests that retrieval-augmented generation is not limited to single-agent settings. It can also support more collaborative forms of intelligence, although this shift introduces additional coordination complexity. In real-world agricultural settings, the framework supports collaboration among agents in environments where conditions are often unstable, and decisions need to be made quickly. A decentralized knowledge graph helps here by enabling agents to learn from each other without depending on a central system.

From an application standpoint, this same design can be applied to areas like autonomous vehicle coordination, smart farming, and other distributed applications where trust, data sharing, and adaptability play a vital role.

Conclusion:

This paper presented a decentralized framework for autonomous agricultural vehicles that combines retrieval-augmented reasoning with blockchain-based trust management. As opposed to conventional approaches that rely on centralized controls or local decision-making, the proposed solution enables autonomous agents to collectively learn from each other's past experiences while preserving data integrity and accountability. The novelty of this work lies in its tightly integrated trust-aware knowledge retrieval mechanism with decentralized execution that allows vehicles to reason not only from their own observations but also from verified experiences contributed by other agents. Through our simulation-based experiments, we showed that the proposed approach improves task success rates and reduces error propagation when compared with local-only, centralized, and unverified decentralized baselines. While the decision latency increases slightly due to retrieval and validation steps, this overhead remains within acceptable limits and is outweighed by the gains in robustness and coordination quality. The results suggest that trusted knowledge sharing becomes increasingly beneficial as the number of autonomous agents grows. The system has been evaluated under controlled simulation settings. Ongoing work is now shifting toward testing the framework in real agricultural environments. These settings vary in soil types, crop patterns, and environmental conditions across different fields. In real deployments, introduces practical constraints such as communication delays, sensor noise, and hardware limitations that are not fully reflected in simulation. As a result, system behavior becomes more variable and less predictable under these sorts of conditions.

Future Work:

Future work will focus on validating the proposed framework using real-world agricultural datasets, improving retrieval mechanisms through advanced embeddings, and optimizing consensus strategies for large-scale deployments. We will explore the adaptive trust models where the trust is not only based on historical validations but also on contextual relevance of the experiences. Furthermore, we plan to evaluate more advanced retrieval strategies and model variants to reduce inference overhead and latency and improve generalization. Extending the framework to support heterogeneous fleets, where vehicles have different capabilities and roles, is another important direction.. Overall, this study demonstrates that combining decentralized intelligence with retrieval-based reasoning and blockchain-based trust is a promising direction for scalable and reliable autonomous farming systems, especially as agricultural automation moves toward larger and more complex deployments.

Acknowledgement: The authors would like to acknowledge the support and facilities provided by the respective institutions during our research.

Author's Contribution: The first author was responsible for implementing the proposed framework and drafting the initial revision. The second author provided the supervision before and during the idea formation and validation of the proposed framework pipeline. The third author was responsible for overseeing the blockchain aspects of the proposed framework. The fourth author was responsible for verifying and validating the proposed framework.

Conflict of interest: The authors declare that there is no conflict of interest of interest in publishing this manuscript in IJIST.

Project details: This research study was not conducted as part of any funded project.

References:

[1] Qingchao Liu, Ruohan Yu, "Autonomous Driving in Agricultural Machinery: Advancing the Frontier of Precision Agriculture," *Actuators*, vol. 14, no. 9, p. 464, 2025, doi: https://doi.org/10.3390/act14090464.

[2] Dehua Zhang, Qingsong Yuan, Lei Meng, Ruixue Xia, Wei Liu & Chunbin Qin, "Reinforcement learning for single-agent to multi-agent systems: from basic theory to industrial application progress, a survey," *Artif. Intell. Rev.*, vol. 59, no. 46, 2026, [Online]. Available: https://link.springer.com/article/10.1007/s10462-025-11439-9

[3] Z. Yang, N. Ma, and Y. Yao, "A Survey on Cooperative Control of Multi-Agent Systems," *Proc. - 2023 2nd Int. Conf. Artif. Intell. Human-Computer Interact. Robot. AIHCIR 2023*, pp. 511–515, 2023, doi: 10.1109/AIHCIR61661.2023.00090.

[4] D. K. Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, "Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks," *arXiv:2005.11401*, 2021, doi: https://doi.org/10.48550/arXiv.2005.11401.
[5] Myeonghwa Lee, Seonho An, Min-Soo Kim, "PlanRAG: A Plan-then-Retrieval Augmented Generation for Generative Large Language Models as Decision Makers," *arXiv:2406.12430*, 2024, [Online]. Available: https://arxiv.org/abs/2406.12430
[6] Andrew Brown, Muhammad Roman, Barry Devereux, "A Systematic Literature Review of Retrieval-Augmented Generation: Techniques, Metrics, and Challenges," *arXiv:2508.06401*, 2025, [Online]. Available: https://arxiv.org/abs/2508.06401
[7] Md Monjurul Karim, Dong Hoang Van, "AI Agents Meet Blockchain: A Survey on Secure and Scalable Collaboration for Multi-Agents," *Futur. Internet*, vol. 17, no. 2, p. 57, 2025, doi: https://doi.org/10.3390/fi17020057.
[8] J. Shen, H. Sheng, S. Wang, R. Cong, D. Yang, and Y. Zhang, "Blockchain-Based Distributed Multiagent Reinforcement Learning for Collaborative Multiobject Tracking Framework," *IEEE Trans. Comput.*, vol. 73, no. 3, pp. 778–788, Mar. 2024, doi: 10.1109/TC.2023.3343102.
[9] "(PDF) Multi-Agent Coordination Strategies vs. Retrieval-Augmented Generation in LLMs: A Comparative Evaluation." Accessed: May 05, 2026. [Online]. Available: https://www.researchgate.net/publication/398597083_Multi-Agent_Coordination_Strategies_vs_Retrieval-Augmented_Generation_in_LLMs_A_Comparative_Evaluation
[10] "Multi-Agent RAG Framework for Entity Resolution: Advancing Beyond Single-LLM Approaches with Specialized Agent Coordination." Accessed: May 06, 2026. [Online]. Available: https://www.mdpi.com/2073-431X/14/12/525
[11] "Is Blockchain the Future of AI Alignment? Developing a Framework and a Research Agenda Based on a Systematic Literature Review." Accessed: May 06, 2026. [Online]. Available: https://www.mdpi.com/2624-800X/5/3/50
[12] "Retrieval-Augmented Generation to Generate Knowledge Assets and Creation of Action Drivers." Accessed: May 06, 2026. [Online]. Available: https://www.researchgate.net/publication/397946484_Generative_AI_for_Sustainable_Smart_Environments_A_Review_of_Energy_Systems_Buildings_and_User-Centric_Decision-Making
[13] Muhammad Murad, Muhammad Ahmed, "Agentic AI Framework to Automate Traditional Farming for Smart Agriculture," *AgriEngineering*, vol. 8, no. 1, p. 8, 2026, doi: https://doi.org/10.3390/agriengineering8010008.
[14] Z. Raza, I. U. Haq and M. Muneeb, "Agri-4-All: A Framework for Blockchain Based Agricultural Food Supply Chains in the Era of Fourth Industrial Revolution," *IEEE Access*, vol. 11, pp. 29851–29867, 2023, doi: 10.1109/ACCESS.2023.3259962.
[15] Mohamad Sheikho Al Jasem, Trevor De Clark, "Toward Decentralized Intelligence: A Systematic Literature Review of Blockchain-Enabled AI Systems," *Information*, vol. 16, no. 9, p. 765, 2025, doi: https://doi.org/10.3390/info16090765.
[16] "Technological Convergence of Blockchain and Artificial Intelligence: A Review and Challenges." Accessed: May 06, 2026. [Online]. Available: https://www.mdpi.com/2079-9292/14/1/84
[17] "Generative AI-Driven Smart Contract Optimization for Secure and Scalable Smart City Services." Accessed: May 06, 2026. [Online]. Available: https://www.mdpi.com/2624-6511/8/4/118
[18] Lavanya Ratnabala, Aleksey Fedoseev, Robinroy Peter, Dzmitry Tsetserukou, "MAGNNET: Multi-Agent Graph Neural Network-based Efficient Task Allocation for Autonomous Vehicles with Deep Reinforcement Learning," *arXiv:2502.02311*, 2025, [Online]. Available: https://arxiv.org/abs/2502.02311

[19] Yingxuan Yang, Huacan Chai, Shuai Shao, Yuanyi Song, Siyuan Qi, Renting Rui, Weinan Zhang, "AgentNet: Decentralized Evolutionary Coordination for LLM-based Multi-Agent Systems," *arXiv:2504.00587*, 2025, [Online]. Available: https://arxiv.org/abs/2504.00587
[20] "GRAG: Graph Retrieval-Augmented Generation - ACL Anthology." Accessed: May 06, 2026. [Online]. Available: https://aclanthology.org/2025.findings-naacl.232/
[21] M. T. Antony James, "Retrieval-Augmented Generation to Generate Knowledge Assets and Creation of Action Drivers," *Appl. Sci.*, vol. 15, no. 11, p. 6247, 2025, doi: https://doi.org/10.3390/app15116247.
[22] Eugene Vorobiov, Ammar Jaleel Mahmood, Salim Rezvani, Robin Chhabra, "ARRC: Advanced Reasoning Robot Control - Knowledge-Driven Autonomous Manipulation Using Retrieval-Augmented Generation," *arXiv:2510.05547*, 2025, [Online]. Available: http://arxiv.org/abs/2510.05547
[23] Yayu Long, Kewei Chen, Long Jin, Mingsheng Shang, "DRAE: Dynamic Retrieval-Augmented Expert Networks for Lifelong Learning and Task Adaptation in Robotics," *arXiv:2507.04661*, 2025, [Online]. Available: https://arxiv.org/abs/2507.04661
[24] Imran Kabir, Md Alimoor Reza, Syed Billah, "Logic-RAG: Augmenting Large Multimodal Models with Visual-Spatial Knowledge for Road Scene Understanding," *arXiv:2503.12663*, 2025, [Online]. Available: https://arxiv.org/abs/2503.12663
[25] Quanting Xie, So Yeon Min, Pengliang Ji, Yue Yang, Tianyi Zhang, Kedi Xu, "Embodied-RAG: General Non-parametric Embodied Memory for Retrieval and Generation," *arXiv:2409.18313*, 2024, [Online]. Available: https://arxiv.org/abs/2409.18313
[26] Ketema Adere, Seble Hailu, Mesfin Abebe, Ararso Kemal, Semere W. Silassie, "A systematic review of artificial intelligence applications in Internet of Things, blockchain, and quantum cryptography," *Discov. Artif. Intell.*, vol. 6, no. 270, 2026, [Online]. Available: https://link.springer.com/article/10.1007/s44163-026-01029-1
[27] Abhirup Khanna, Sapna Jain, "Generative AI and Blockchain-Integrated Multi-Agent Framework for Resilient and Sustainable Fruit Cold-Chain Logistics," *Foods*, vol. 14, no. 17, p. 3004, 2025, doi: https://doi.org/10.3390/foods14173004.
[28] Awid Vaziry, Sandro Rodriguez Garzon, Axel Küpper, "Towards Multi-Agent Economies: Enhancing the A2A Protocol with Ledger-Anchored Identities and x402 Micropayments for AI Agents," *arXiv:2507.19550*, 2025, [Online]. Available: https://arxiv.org/abs/2507.19550
[29] "What is Ollama? A Beginner's Guide To This Platform." Accessed: May 06, 2026. [Online]. Available: https://blog.udemy.com/what-is-ollama/
[30] G. L. Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, "LLaMA: Open and Efficient Foundation Language Models," *arXiv:2302.13971*, 2023, doi: https://doi.org/10.48550/arXiv.2302.13971.
[31] Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, "Alpaca: A Strong, Replicable Instruction-Following Model," *Stanford Univ. Press*, 2023, [Online]. Available: https://crfm.stanford.edu/2023/03/13/alpaca.html
[32] Philip Sun, David Simcha, Dave Dopson, Ruiqi Guo, Sanjiv Kumar, "SOAR: Improved Indexing for Approximate Nearest Neighbor Search," *arXiv:2404.00774*, 2024, [Online]. Available: https://arxiv.org/abs/2404.00774