ARES-SC:
Blockchain for Multi-Agent Learning

Masashi Sugimoto
http://orcid.org/0000-0001-6201-1475
Tomakomai College, Japan

ABSTRACT

Multi-agent systems increasingly require transparency alongside performance optimization. This chapter presents ARES-SC, integrating blockchain with multi-agent reinforcement learning to address trust-performance trade-offs. The system combines (1) adaptive resolution for dynamic resource allocation, (2) semantic communication for interpretable coordination, and (3) blockchain for immutable decision recording. Experimental validation across three scenarios reveals important trade-offs. ARES-CORE achieved highest performance (normalized scores 2900-3500) but with elevated overhead (6657-6875%). ARES-BC provided competitive performance (2500-3000) with complete cryptographic verification at measurable costs (625-737 blockchain units). All methods exhibited significant computational overhead (6200-6800%). This work provides honest assessment of when transparency features justify computational costs, offering practical guidance for trustworthy multi-agent systems.

1. INTRODUCTION

1.1 The Challenge of Trustworthy Multi-Agent Systems

Imagine deploying autonomous agents in a hospital emergency department. While these systems excel at optimizing resource allocation in controlled environments, real-world deployment faces a fundamental challenge: stakeholders need to understand and verify AI decisions that affect human welfare. When a medical review board asks “Why did the system prioritize this patient?” or “Can you prove the allocation followed approved protocols?”, raw performance metrics become insufficient.

The artificial intelligence field has reached a critical juncture where technical capability outpaces social acceptance. Multi-agent systems demonstrate impressive coordination in simulated environments, yet their deployment in healthcare, transportation, and finance faces resistance due to their “black box” nature (Amodei et al., 2016; Gunning & Aha, 2019). This resistance reflects not just technical limitations but a fundamental mismatch between optimization-focused algorithms and human requirements for transparency and accountability.

DOI: 10.4018/406042

Traditional reinforcement learning approaches prioritize performance metrics —reward maximization, convergence speed, computational efficiency. However, real-world deployment requires additional considerations: regulatory compliance, multi-stakeholder trust, decision auditability, and failure accountability (Guidotti et al., 2018). The central question becomes: how do we design systems that balance high performance with the transparency demands of human society?

1.2 From Adaptive Resolution to Transparent Coordination

The Adaptive Resolution Exploration Strategies for Competitive Optimization of Reinforcement Learning (ARES-CORE) framework addresses computational efficiency by dynamically adjusting learning granularity based on environmental complexity (Saad-Falcon et al., 2023). Like human attention systems that focus intensely on complex situations while handling routine tasks automatically, ARES-CORE allocates computational resources where they provide the most value.

However, efficiency improvements alone proved insufficient for addressing trust requirements. This limitation motivated the development of ARES-SC (ARES with Semantic Communication), which enables agents to exchange semantically meaningful information using natural language processing models like BERT (Devlin et al., 2018). Instead of sharing opaque numerical vectors, agents can communicate interpretable concepts that human supervisors can understand and verify.

Yet even interpretable communication fails to fully address accountability requirements in critical applications. Medical ethics boards, financial regulators, and safety authorities need more than understandable decisions—they need verifiable ones with immutable audit trails.

1.3 Blockchain Integration: Beyond Cryptocurrency Applications

Blockchain technology offers a practical solution for creating tamper-resistant decision records. Beyond its origins in cryptocurrency, blockchain provides immutable, distributed ledgers that can record every agent action, communication, and policy update with cryptographic verification (Nakamoto, 2008; Tapscott & Tapscott, 2016).

For multi-agent systems, blockchain integration serves three critical functions:

- **Decision Auditability**: Every action and its justification are permanently recorded
- **Multi-Stakeholder Coordination**: Organizations can coordinate without trusting a central authority
- **Regulatory Compliance**: Automated generation of audit trails for regulatory review

The blockchain component is not theoretical—it addresses real deployment barriers. When regulatory bodies require proof that AI systems followed approved protocols, blockchain records provide cryptographically verifiable evidence of compliance.

1.4 Research Scope and Honest Assessment

This research makes specific contributions while acknowledging clear limitations:

Primary Contributions:

Related Content