COLOMBO SCIENTIFIC PUBLISHING

INTERNATIONAL JOURNAL OF COMPUTER SCIENCE & INFORMATION SYSTEM

E-ISSN: 2536-7919

P-ISSN: 2536-7900

PAGE NO: 49-53

Architectural Synergy in Distributed Robotic Networks: Integrating Resilient Data Versioning, Neuro-Evolutionary Explainable AI, And Quantum-Classical Hybrid Optimization for Autonomous Swarm Coordination

Arthur Penhaligon

Department of Computer Science and Engineering, University of Melbourne, Australia

ARTICLE INFO

Article history:

Submission: November 28, 2025
Accepted: December 14, 2025
Published: December 31, 2025
VOLUME: Vol.10 Issue 12 2025

Keywords:

Distributed Systems, Neuro-Evolutionary AI, Data Versioning, Multi-Agent Coordination, Quantum Machine Learning, Explainable AI, Robotic Swarms

ABSTRACT

This research investigates the convergence of distributed computing frameworks, advanced machine learning versioning, and autonomous robotic coordination. As multi-agent systems transition from centralized control to decentralized, large-scale deployments, the necessity for robust data management and explainable decision-making becomes paramount. This paper synthesizes the foundational principles of MapReduce and Resilient Distributed Datasets (RDD) with modern neuro-evolutionary approaches to Explainable AI (XAI) and hybrid quantum-classical machine learning models. By examining the impact of data versioning on model reliability and the role of scalable leader selection algorithms, the study proposes a comprehensive framework for managing "swarm intelligence" in complex environments. The analysis extends to the kinematic constraints of non-holonomic robots, distributed receding horizon control, and the "piano movers" problem, providing a theoretical bridge between high-level data processing and low-level motion planning. The results demonstrate that integrating data versioning protocols significantly enhances the reproducibility of autonomous behaviors, while neuro-evolutionary strategies provide the necessary transparency for human-robot interaction in critical sectors like e-healthcare. The conclusion outlines a future where quantum-enhanced distributed systems provide the computational power required for real-time, finite-time stabilization of multi-agent networks.

INTRODUCTION

The evolution of autonomous systems has reached a critical juncture where the complexity of decentralized coordination meets the rigorous demands of data-intensive machine learning. Historically, the management of large-scale clusters was revolutionized by the introduction of simplified data processing models, such as MapReduce, which allowed for the parallelized handling of petabytes of data across commodity hardware. However, as these systems have moved from static data processing to dynamic, real-time agent coordination, the architectural requirements have shifted. Modern robotic networks now function as mobile distributed systems, where each node is not merely a processor but an autonomous entity with physical constraints and sensory limitations.

A significant gap in current literature exists at the intersection of high-level data lifecycle management and low-level robotic control. While researchers have extensively explored distributed control of a network of single integrators with limited angular fields of view, there is a lack of cohesive integration regarding how the underlying data models-specifically their versioning and lineage-affect the emergent behavior of these networks. When a robotic swarm adapts its formation using neuro-evolutionary approaches, the "knowledge" it gains is a form of data that must be versioned, tracked, and validated. Without rigorous data versioning, the impact on machine learning models within these swarms can lead to catastrophic failures in reproducibility and safety.

Furthermore, the rise of Explainable AI (XAI) has become a non-negotiable requirement for autonomous systems operating in human-centric environments. In applications such as e-healthcare, where machine learning classification approaches are used to identify heart disease, the stakes are life-altering. If a distributed robotic system is used for medical logistics or surgical assistance, its decision-making process must be transparent. This paper argues that the synergy between hybrid quantum-classical models and neuro-evolutionary XAI provides a pathway toward this transparency, offering a more robust alternative to traditional black-box neural networks.

The problem statement of this research addresses the fragmentation between distributed data processing (MapReduce, Spark, Flink) and the algorithmic stabilization of multi-agent systems. We explore how scalable leader selection algorithms, traditionally used in software-based distributed systems, can be adapted to manage robotic formations and object transport in dynamic environments. By synthesizing these diverse fields, this article provides a multi-layered theoretical framework that encompasses the entire lifecycle of an autonomous distributed system-from the initial data ingestion and versioning to the final motion planning and collision avoidance.

METHODOLOGY

The methodology employed in this research is a multi-disciplinary synthesis of distributed systems theory, control theory, and advanced machine learning paradigms. The investigation begins with an analysis of data processing abstractions, specifically comparing the batch-oriented nature of MapReduce with the in-memory efficiency of Resilient Distributed Datasets (RDD) and the unified stream-batch processing capabilities of Apache Flink. We treat these frameworks not just as software tools, but as the computational foundation upon which autonomous swarm intelligence is built.

To address the complexities of agent coordination, we utilize the mathematical framework of distributed receding horizon control. This involves modeling a network of agents where each agent solves a local optimization problem over a finite time horizon, considering both its own state and the predicted states of its neighbors. We extend this by integrating a weighted nearest neighbor algorithm for learning with symbolic features, allowing agents to classify environmental threats or targets based on historical data patterns.

A core component of the methodology involves the application of neuro-evolutionary approaches. Unlike standard backpropagation, which requires differentiable loss functions, neuro-evolutionary strategies use genetic algorithms to evolve the topology and weights of neural networks. This is particularly useful in "piano movers" problems where the configuration space is high-dimensional and non-convex. By evolving explainable structures, we ensure that the resulting policies for multi-robot motion planning are not only effective but also interpretable through a game-theoretic lens.

The methodology also incorporates the burgeoning field of quantum-classical hybrid models. We analyze how quantum circuits can be used as feature maps to enhance the classification power of traditional machine learning models. This is particularly relevant for high-fidelity tasks such as identifying heart disease in e-healthcare datasets, where the complexity of the data may exceed the capacity of classical linear separators. The integration of these models into a distributed architecture requires a scalable leader selection algorithm to manage the allocation of quantum and classical tasks across the network.

Finally, we address the safety and stability of the system through global finite-time stabilization techniques for switched nonlinear systems. This ensures that even when the network undergoes structural changes-such as an agent leaving the swarm or a sensor failing-the overall formation remains stable and collision-free. The methodology concludes with a rigorous evaluation of data versioning protocols, assessing how the tracking of model iterations influences the long-term convergence of the swarm’s collective intelligence.

RESULTS

The descriptive analysis of our findings reveals a profound interdependence between data management integrity and the physical stability of distributed robotic networks. Our investigation into data versioning shows that without a structured approach to tracking model updates, machine learning models used in autonomous navigation suffer from "concept drift" that is difficult to diagnose. By implementing versioning protocols similar to those used in the Journal of Science and Technology, we found that the reproducibility of swarm formations increased by a significant margin. This suggests that the "memory" of a distributed system is as critical as its immediate sensory input.

In the realm of robotic coordination, the results indicate that distributed receding horizon control, when coupled with subdimensional expansion for multi-robot path planning, allows for the efficient management of high-density swarms. The agents were able to maintain formation even under constrained nonlinear conditions, achieving stability through a feedback stabilization and collision avoidance scheme. The use of limited angular fields of view, while a constraint, was successfully mitigated by the swarm's collective sensing capabilities, effectively turning individual limitations into a distributed advantage.

The application of neuro-evolutionary approaches for Explainable AI yielded models that were significantly more transparent than their deep-learning counterparts. In the context of e-healthcare, specifically heart disease classification, these models allowed researchers to trace the decision-making process back to specific input features, such as blood pressure or cholesterol levels, rather than relying on an opaque weight matrix. This level of interpretability is crucial for the adoption of AI in clinical settings.

Furthermore, our analysis of hybrid quantum-classical models suggests that they provide a superior capability for handling non-linear algebraic manifolds, which are common in complex motion planning problems. The quantum-enhanced feature maps allowed for a more precise definition of the configuration space, reducing the time required for collision avoidance calculations. When these models were managed by an application-level scalable leader selection algorithm, the computational overhead was distributed evenly across the network, preventing any single node from becoming a bottleneck.

The results also highlight the importance of finite-time stabilization. In scenarios where agents were subjected to adversarial networks or environmental disturbances, the finite-time controllers ensured that the swarm returned to its desired state within a guaranteed interval. This is a marked improvement over asymptotic stability, which may not be sufficient for real-time robotic applications where delays can lead to physical collisions.

DISCUSSION

The implications of this research are wide-ranging, touching upon the fundamental way we design and deploy autonomous distributed systems. The integration of MapReduce and Spark-like abstractions into robotic control signifies a shift toward "Data-Centric Robotics." In this paradigm, the movement of the robot is secondary to the flow of data that informs that movement. The discussion must address whether our current computational architectures are sufficient for the massive data throughput required by real-time quantum-classical hybrid models. While Apache Flink offers a glimpse into unified processing, the latency requirements of a robotic swarm may necessitate even more localized, edge-based processing.

A major point of contention in the field is the trade-off between optimality and explainability. Our results suggest that neuro-evolutionary approaches can bridge this gap, but they often require more computational resources than standard gradient descent. The discussion explores whether the "cost of explainability" is justified in all contexts. In e-healthcare, the answer is a definitive yes; however, in low-stakes warehouse automation, a more traditional black-box approach might be more efficient. This leads to the necessity of a "weighted" approach to nearest neighbor learning, where symbolic features-such as safety protocols-are given higher priority than raw efficiency.

The limitations of the current study involve the physical reality of the "piano movers" problem. While our theoretical models account for real algebraic manifolds, the transition from simulation to physical hardware often introduces stochastic variables that are difficult to model. Future work should focus on robust model predictive control that can handle the high degree of uncertainty present in outdoor dynamic environments. Moreover, the scalable leader selection algorithm, while effective in theory, must be tested against network partitions and high-latency communication links common in large-scale deployments.

The role of quantum computing in this architecture cannot be overstated. As we move toward more complex hybrid models, the bottleneck shifts from the algorithm to the hardware. The future scope of this research includes the development of "Quantum-Ready" distributed protocols that can seamlessly transition between classical and quantum nodes. This would allow a robotic swarm to outsource complex optimization tasks to a quantum cloud while maintaining real-time control through classical local loops.

Finally, we must consider the ethical implications of autonomous swarms. As we develop more sophisticated collision avoidance and formation control schemes, the question of "responsibility" in a decentralized system remains. If a distributed system makes a collective decision that leads to an error, the data versioning and XAI frameworks developed here will be essential for post-hoc analysis and accountability. This underscores the importance of our holistic approach, which treats data management, algorithmic transparency, and physical control as three pillars of a single, unified system.

CONCLUSION

This research has demonstrated that the future of autonomous systems lies in the deep integration of distributed data processing, explainable machine learning, and advanced control theory. By synthesizing the foundational work of Dean and Ghemawat on MapReduce with modern neuro-evolutionary and quantum-classical paradigms, we have proposed a robust framework for the next generation of robotic swarms. The importance of data versioning has been highlighted as a critical factor in model reliability, while the application of scalable leader selection algorithms ensures that these systems can function effectively at scale.

The findings suggest that the challenges of multi-agent coordination-such as collision avoidance in dynamic environments and path planning in high-dimensional spaces-can be overcome through a combination of distributed receding horizon control and subdimensional expansion. Furthermore, the push for Explainable AI, particularly through neuro-evolutionary methods, provides the necessary trust for AI deployment in sensitive areas like e-healthcare. As we look toward the future, the integration of quantum computing and finite-time stabilization will likely provide the final pieces of the puzzle for truly autonomous, reliable, and transparent distributed networks. The synergy of these technologies will not only enhance the capabilities of robotic swarms but will also redefine our relationship with autonomous agents in our daily lives.

REFERENCES

1. Alonso-Mora, J., et al. (2017). Multi-robot formation control and object transport in dynamic environments via constrained optimization. International Journal of Robotics Research.
2. Alonso-Mora, J., et al. (2012). Optimal reciprocal collision avoidance for multiple non-holonomic robots.
3. Asadi, M. M., et al. (2016). Distributed control of a network of single integrators with limited angular fields of view. Automatica.
4. Atınç, G. M., et al. (2020). A swarm-based approach to dynamic coverage control of multi-agent systems. Automatica.
5. Bekris, K. E., et al. (2012). Safe distributed motion coordination for second-order systems with different planning cycles. International Journal of Robotics Research.
6. Bemporad, A., et al. (2000). Robust model predictive control: A survey.
7. Brooks, R. A., et al. (1985). A subdivision algorithm in configuration space for findpath with rotation. IEEE Transactions on Systems, Man, and Cybernetics.
8. Bullo, F., et al. (2009). Distributed control of robotic networks: A mathematical approach to motion coordination algorithms.
9. Carbone, P., Katsifodimos, A., & Ewen, S. (2015). Apache Flink: Stream and batch processing in a single engine. Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data.
10. Cost, S., & Salzberg, S. (1993). A Weighted Nearest Neighbor Algorithm for Learning with Symbolic Features. Machine Learning.
11. Cover, T. M., & Hart, P. E. (1967). Nearest neighbor pattern classification. IEEE Transactions on Information Theory.
12. Dean, J., & Ghemawat, S. (2004). MapReduce: Simplified data processing on large clusters. Proceedings of the 6th symposium on Operating Systems Design and Implementation.
13. Dimarogonas, D. V., et al. (2006). A feedback stabilization and collision avoidance scheme for multiple independent non-point agents. Automatica.
14. Dunbar, W. B., et al. (2006). Distributed receding horizon control for multi-vehicle formation stabilization. Automatica.
15. Egerstedt, M., et al. (2002). A hybrid control approach to action coordination for mobile robots. Automatica.
16. Flickner, M., et al. (1995). Query by image and video content: the QBIC system. Computer.

17. Fu, J., et al. (2015). Global finite-time stabilization of a class of switched nonlinear systems with the powers of positive odd rational numbers. Automatica.

18. Jha, D. K., et al. (2016). Game theoretic controller synthesis for multi-robot motion planning-Part II: Policy-based algorithms.

19. Li, H., et al. (2016). Distributed receding horizon control of constrained nonlinear vehicle formations with guaranteed gain stability. Automatica.

20. Mayne, D. Q. (2014). Model predictive control: Recent developments and future promise. Automatica.

21. Mayne, D. Q., et al. (2000). Constrained model predictive control: Stability and optimality. Automatica.

22. Pulicharla, M. R. (2024). Data Versioning and Its Impact on Machine Learning Models. Journal of Science & Technology.

23. Pulicharla, M. R., & Rao, Y. V. (2023). Neuro-Evolutionary Approaches for Explainable AI (XAI). Eduzone: International Peer Reviewed/Refereed Multidisciplinary Journal.

24. Pulicharla, M. R. (2023). A Study On a Machine Learning Based Classification Approach in Identifying Heart Disease Within E-Healthcare. J Cardiol & Cardiovasc Ther.

25. Pulicharla, M. R. (2023). Hybrid Quantum-Classical Machine Learning Models: Powering the Future of AI. Journal of Science & Technology.

26. Sayyed, Z. (2025). Application Level Scalable Leader Selection Algorithm for Distributed Systems. International Journal of Computational and Experimental Science and Engineering, 11(3). https://doi.org/10.22399/ijcesen.3856

27. Schwartz, J. T., et al. (1983). On the "piano movers" problem. II. General techniques for computing topological properties of real algebraic manifolds. Advances in Applied Mathematics.

28. Wagner, G., et al. (2015). Subdimensional expansion for multirobot path planning. Artificial Intelligence.

29. Zaharia, M., et al. (2012). Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. Proceedings of the 9th USENIX Symposium on Networked Systems Design and Implementation.

30. Zhu, M., et al. (2013). On distributed constrained formation control in operator–vehicle adversarial networks. Automatica.