Quantum-Enhanced Learning and Control for Multi-agent Systems

Doctoral Consortium

Meet Gandhi
Colorado School of Mines
Golden, Colorado, USA
meet_gandhi@mines.edu

ABSTRACT

With multi-agent systems advancing to high-dimensional and uncertain spaces, classical learning and control methods encounter fundamental challenges related to the curse of dimensionality and limited expressivity. Quantum computing offers the potential to overcome the shortcomings by embedding data into exponentially large Hilbert spaces, capturing complex correlations. We first propose a Distributed Quantum Gaussian Process (DQGP) framework enabling agents to collaboratively learn a high-fidelity global model of the environment through improved modeling capabilities and scalability. Numerical evaluations on non-stationary NASA SRTM datasets demonstrate the enhanced predictive and uncertainty estimation performance of DQGP compared to the classical Distributed Gaussian Processes. The findings lead to the next research phase: developing a Quantum-enhanced Learning Model Predictive Control architecture that results in robust, adaptive, and scalable coordination and control of multiple agents in complex scenarios.

KEYWORDS

Quantum Computing; Gaussian Processes; Learning-based Control; Multi-agent Systems

ACM Reference Format:

Meet Gandhi. 2026. Quantum-Enhanced Learning and Control for Multi-agent Systems: Doctoral Consortium. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/KNGI7618

1 INTRODUCTION

For large-scale multi-agent systems, coordination and control present a fundamental challenge, especially when the system needs to learn complex models of the environment and perform optimal control under uncertainty. Classical models for learning and distributed control provide a robust framework with theoretical guarantees. However, an increase in the number of agents or the dimensionality of the state-space exposes the scalability challenges of established approaches. The surge in dimensionality introduces a level of complexity that exceeds the expressive power of the classical frameworks. This research proposes a novel combination of quantum information theory with cooperative learning-based control, exploring the efficacy of quantum-enhanced algorithms in mitigating the classical constraints. By leveraging the high expressivity of quantum kernels using the Hilbert latent space representation and the computational speedup of quantum-enhanced optimization schemes; the aim is enabling multiple agents to operate with high fidelity in large-scale and uncertain environmental spaces.

Attaining high-fidelity control requires a modeling methodology that can accurately characterize the uncertain residual dynamics in a computationally tractable manner. Gaussian processes (GPs), as an inherently probabilistic learning technique, satisfy the need for reliability through accurate predictions and principled uncertainty estimation; however, training a GP that models the intrinsic dynamics of an unknown environment using N samples involves $O(N^3)$ computations and $O(N^2)$ memory. Distributed GP (DGP) methods [4, 6, 15, 18, 23] effectively overcome the scalability limitations of standard GP by distributing both data storage and computational effort among multiple agents. Distributed Gaussian Processes employ kernel functions [14] to model the correlations among the data points by projecting them into a high-dimensional feature space. The mapping enables DGPs to capture complex relationships. However, the classical kernels are limited in their expressivity due to the underlying mathematical formulation that is tractable on classical hardware.

Motivated by the need for higher-dimensional representation, the first phase of the research proposes a Distributed Quantum Gaussian Process (DQGP) [9] framework for accurately learning the environment and unknown dynamics in multi-agent systems. The emerging field of quantum computing [2, 3, 19] is leveraged to model high-dimensional correlations. A major advancement in quantum computing in recent years has been the introduction of quantum kernel functions [12], which have led to the development of Quantum Gaussian Processes (QGPs) [20] among others [1, 21]. Quantum kernels incorporate parameterized quantum circuits, called Quantum Encoding Circuits [22], to embed classical data into the quantum domain. Considerable research has focused on designing kernels to achieve optimal alignment between quantum feature spaces and classical data labels. Beyond QGPs, the quantum kernels have proven beneficial in other domains of quantum machine learning tasks [1, 21]. Accordingly, our proposed DQGP approach yields up to a two-fold reduction in prediction error by leveraging the enhanced expressivity of quantum kernels while maintaining computational scalability. The success of the presented learning framework establishes a crucial foundation for the subsequent control research, where the objective is to combine learning in the Quantum Hilbert space with quantum-enhanced model predictive control architectures for multi-agent systems.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/KNGI7618

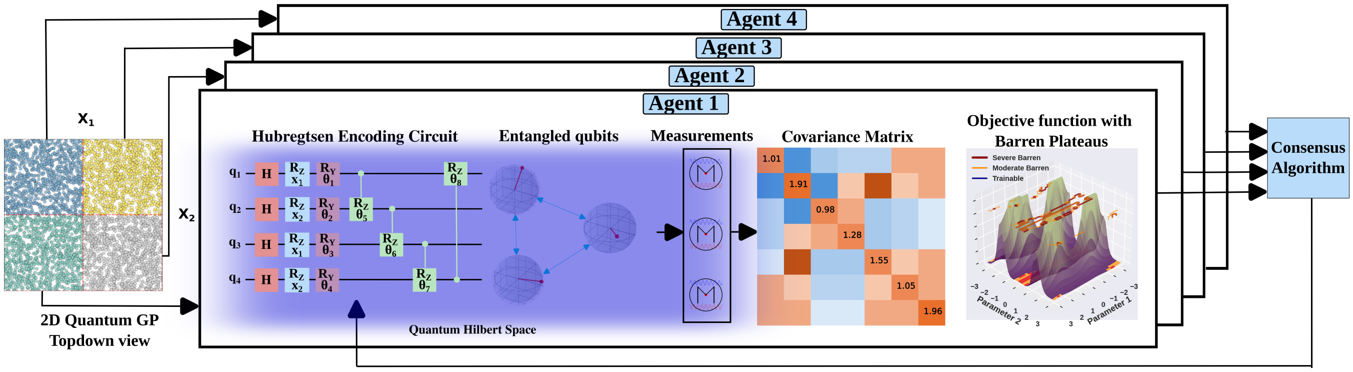

Figure 1: The structure of the proposed DQGP with 4 agents. The consensus algorithm is the proposed DR-ADMM optimization.

2 DISTRIBUTED QUANTUM GAUSSIAN PROCESSES

The goal in this preliminary work is to perform scalable and highly expressive distributed probabilistic modeling using quantum kernels and a novel Riemannian optimization framework. The quantum kernel lifts classical data into a high-dimensional Hilbert space via quantum encoding circuits. Fig. 1 shows the Hubregtsen Encoding Circuit [13] along with the proposed methods: DQGP and DR-ADMM. The Hubregtsen Encoding Circuit induces quantum effects by utilizing the Hadamard gates H for *superposition* and conditional-R_Z gates to generate *entanglement*, while parameterized R_Z and R_Y rotational gates encode the classical data. This results in the modeling of correlations that classical kernels cannot represent efficiently. Moreover, the Distributed Riemannian ADMM (DR-ADMM) consensus algorithm is specifically designed to optimize quantum hyperparameters, corresponding to qubit rotation angles, on a toroidal manifold rather than Euclidean space. Classical GP optimization would treat these angles as unconstrained real variables and fail to capture the intrinsic periodic geometry. While the work focuses on DQGP regression, DR-ADMM provides a distributed optimization foundation for applications involving quantum circuit parameters.

We evaluate the performance of DQGP through numerical experiments conducted on NASA's Shuttle Radar Topography Mission (SRTM) [8]. The non-stationary [5] property of the dataset is useful to assess whether the learned model can fully adapt to the local features, thus demonstrating its suitability for solving complex problems. Negative log predictive density (NLPD) and normalized RMSE (NRMSE) metrics provide a comprehensive assessment of the approach's predictive and uncertainty quantification capabilities. Moreover, apx-GP [23] is used as a classical DGP baseline approach for benchmarking. Aggregating results from different dataset sizes $N = \{500, 5000\}$ and number of agents $M = \{4, 8, 27\}$, DQGP achieves $65.2\% \pm 16.1\%$ lower $\text{NRMSE}_{\text{test}}$ than apxGP, and $91.7\% \pm 11.2\%$ improvement in $\text{NLPD}_{\text{test}}$ relative to apxGP. The trend suggests that our approach (DQGP) scales effectively with network size. The results highlight the potential of DQGP for scalable probabilistic modeling on hybrid classical-quantum systems, paving the way for quantum-enhanced autonomous systems.

The demonstrated efficacy of DQGP motivated our research towards assessing the expressive power of quantum kernels. We aim to quantify the modeling advantages gained from using the Quantum Hilbert space instead of the classical embedding space. Ultimately, the study aims to provide insights and broaden the perspective of the quantum-enhanced optimization field, demonstrating that the quantum advantage extends beyond the computational runtime speedup by including the enhanced expressive power of quantum kernels.

3 QUANTUM-ENHANCED MULTI-AGENT CONTROL

The next research trajectory involves bridging the gap between quantum-enhanced learning and model-based control. [11] presents a comprehensive review of learning-based control methodologies and highlights that employing probabilistic learning methods for real-time control in multi-agent systems necessitates overcoming the challenges regarding safety guarantees and computational latency. The issue of safety in GP-based Learning Model Predictive Control (LMPC) for single agents is discussed in [10]. On the other hand, [17] focuses on the scalability issue of GP-based LMPC for multi-agent systems. The application of DGP approaches for LMPC in a multi-agent setting is largely unexplored, presenting a critical area that this research aims to address.

To formulate a baseline approach for LMPC, we explore the robust and adaptive GP-RAMPC framework presented in [7] for a single agent. Then our objective is to scale the approach for multiple agents. We aim to employ centralized and decentralized GP training methodologies like apx-GP, gapx-GP, DEC-apx-GP, and DEC-gapx-GP presented in [15]. For GP prediction, we plan to explore expert aggregation techniques like PoE, NPAE, BCM, rBCM, and their respective decentralized versions as shown in [16]. These approaches provide a good distributed GP learning backdrop for the robust and adaptive GP-RAMPC to be implemented on multi-agent systems. Following a rigorous assessment of the established distributed GP-RAMPC framework, we plan to focus on the limited expressivity and computational scalability challenges by developing a quantum-enhanced distributed GP-RAMPC architecture using quantum kernels.

REFERENCES

[1] Luca Arceci, Viacheslav Kuzmin, and Rick Van Bijnen. 2024. Gaussian process model kernels for noisy optimization in variational quantum algorithms. *arXiv preprint arXiv:2412.13271* (2024).

[2] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A Buell, et al. 2019. Quantum supremacy using a programmable superconducting processor. *Nature* 574, 7779 (2019), 505–510.

[3] Julian Berberich and Daniel Fink. 2024. Quantum computing through the lens of control: A tutorial introduction. *IEEE Control Systems* 44, 6 (2024), 24–49.

[4] Tsung-Hui Chang, Mingyi Hong, and Xiangfeng Wang. 2014. Multi-agent distributed optimization via inexact consensus ADMM. *IEEE Transactions on Signal Processing* 63, 2 (2014), 482–497.

[5] Weizhe Chen, Roni Khardon, and Lantao Liu. 2024. Adaptive robotic information gathering via non-stationary Gaussian processes. *The International Journal of Robotics Research* 43, 4 (2024), 405–436.

[6] Marc Deisenroth and Jun Wei Ng. 2015. Distributed gaussian processes. In *International conference on machine learning*. PMLR, 1481–1490.

[7] Mathieu Dubied, Amon Lahr, Melanie N Zeilinger, and Johannes Köhler. 2025. A robust and adaptive MPC formulation for Gaussian process models. *arXiv preprint arXiv:2507.02098* (2025).

[8] Tom G Farr, Paul A Rosen, Edward Caro, Robert Crippen, Riley Duren, Scott Hensley, Michael Kobrick, Mimi Paller, Ernesto Rodriguez, Ladislav Roth, et al. 2007. The shuttle radar topography mission. *Reviews of geophysics* 45, 2 (2007).

[9] Meet Gandhi and George P Kontoudis. 2026. Distributed Quantum Gaussian Processes for Multi-Agent Systems. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. https://doi.org/10.65109/ADPL7324

[10] Lukas Hewing, Juraj Kabzan, and Melanie N Zeilinger. 2019. Cautious model predictive control using gaussian process regression. *IEEE Transactions on Control Systems Technology* 28, 6 (2019), 2736–2743.

[11] Lukas Hewing, Kim P Wabersich, Marcel Menner, and Melanie N Zeilinger. 2020. Learning-based model predictive control: Toward safe learning in control. *Annual Review of Control, Robotics, and Autonomous Systems* 3, 1 (2020), 269–296.

[12] Hsin-Yuan Huang, Michael Broughton, Masoud Mohseni, Ryan Babbush, Sergio Boixo, Hartmut Neven, and Jarrod R McClean. 2021. Power of data in quantum machine learning. *Nature communications* 12, 1 (2021), 2631.

[13] Thomas Hubregtsen, David Wierichs, Elies Gil-Fuster, Peter-Jan HS Derks, Paul K Faehrmann, and Johannes Jakob Meyer. 2022. Training quantum embedding kernels on near-term quantum computers. *Physical Review A* 106, 4 (2022), 042431.

[14] Motonobu Kanagawa, Philipp Hennig, Dino Sejdinovic, and Bharath K Sriperumbudur. 2018. Gaussian processes and kernel methods: A review on connections and equivalences. *arXiv preprint arXiv:1807.02582* (2018).

[15] George P Kontoudis and Daniel J Stilwell. 2024. Scalable, federated Gaussian process training for decentralized multi-agent systems. *IEEE Access* 12 (2024), 77800–77815.

[16] George P Kontoudis and Daniel J Stilwell. 2025. Multi-Agent Federated Learning using Covariance-Based Nearest Neighbor Gaussian Processes. *IEEE Transactions on Machine Learning in Communications and Networking (TMLCN)* 4 (2025), 115–138. https://doi.org/10.1109/TMLCN.2025.3643409

[17] Viet-Anh Le and Truong X Nghiem. 2020. Gaussian process based distributed model predictive control for multi-agent systems using sequential convex programming and ADMM. In *2020 IEEE Conference on Control Technology and Applications (CCTA)*. IEEE, 31–36.

[18] Haitao Liu, Jianfei Cai, Yi Wang, and Yew Soon Ong. 2018. Generalized Robust Bayesian Committee Machine for Large-scale Gaussian Process Regression. In *International Conference on Machine Learning*. 3131–3140.

[19] John Preskill. 2018. Quantum computing in the NISQ era and beyond. *Quantum* 2 (2018), 79.

[20] Frederic Rapp and Marco Roth. 2024. Quantum gaussian process regression for bayesian optimization. *Quantum Machine Intelligence* 6, 1 (2024), 5.

[21] Alistair WR Smith, AJ Paige, and MS Kim. 2023. Faster variational quantum algorithms with quantum kernel-based surrogate models. *Quantum Science and Technology* 8, 4 (2023), 045016.

[22] Chelsea A Williams, Annie E Paine, Hsin-Yu Wu, Vincent E Elfving, and Oleksandr Kyriienko. 2023. Quantum chebyshev transform: Mapping, embedding, learning and sampling distributions. *arXiv preprint arXiv:2306.17026* (2023).

[23] Ang Xie, Feng Yin, Yue Xu, Bo Ai, Tianshi Chen, and Shuguang Cui. 2019. Distributed Gaussian processes hyperparameter optimization for big data using proximal ADMM. *IEEE Signal Processing Letters* 26, 8 (2019), 1197–1201.