Postdoctoral Research Topic

- **Title of the proposed topic:** Quantum-to-Classical Knowledge Distillation for Robotics: A Quantum Teacher and a Classical Student
- **Research axis:** Axe 3 - AI for Smart and Secure Spaces
- **Supervisor:** Ezio MALIS
- **Research group:** ACENTAURI project-team, Inria Center at Université Côte d'Azur

Research teams

ACENTAURI is a robotic team located in Sophia Antipolis that studies and develop intelligent, autonomous, and mobile robots that collaborate between them to achieve challenging tasks in dynamic environments. The team tackle perception, decision, and control problems for multi-robot collaboration by proposing an original hybrid rule-driven/data driven approach to artificial intelligence and by studying efficient optimization algorithms. The team focus on robotic applications like environment monitoring and transportation of people and goods. In these applications, several robots will share multi-sensor information eventually coming from infrastructure. The effectiveness of the proposed approaches is demonstrated on real robotic systems like cars AGVs and UAVs together with industrial partners.

1 Context

Modern robotic systems increasingly rely on learning-based perception, prediction, and decision-making modules that must operate under strict constraints on latency, compute, energy, and reliability. In many robotic platforms (mobile robots, drones, autonomous vehicles, ...), inference must execute on embedded or edge hardware while remaining robust to sensor noise, distribution shifts, and safety-critical failures. Robotics literature emphasizes that autonomy pipelines integrate learned components with planning and control, and that practical deployment requires careful attention to computational limits and robustness in real environments [1].

Quantum computers are increasingly accessible for experimentation, yet they remain difficult to deploy directly onboard robots due to hardware availability, latency, sampling cost, and noise. Previous work on quantum machine learning (QML) emphasize that these constraints shape what is feasible with noisy intermediate-scale quantum devices and motivate hybrid workflows that use quantum computation selectively during training rather than at deployment [2].

At the same time, QML models (such as quantum kernel methods and variational quantum classifiers) aim to exploit quantum feature spaces and parameterized circuits to represent decision functions that may be difficult to capture with compact classical models [3, 4]. However, even when a quantum model shows promise, running it as a production inference engine is often impractical in robotic systems.

Knowledge distillation provides a principled way to transfer predictive behavior from a teacher model to a smaller student model by training the student to mimic the teacher's outputs (often using soft probability targets), enabling efficient deployment [5]. This proposal investigates *quantum-to-classical* distillation for robotics: a quantum teacher generates supervision signals during training, while the deployed model remains a standard classical student suitable for real-time inference.

2 Postdoc Subject

The main goal of this postdoc is to develop a **quantum-to-classical teacher–student learning framework** in which a **teacher model is implemented on a quantum computer** and a **student model is implemented on a classical computer**, with the aim of achieving deployable classical inference for robotic perception and decision-making that benefits from quantum-derived supervision signals.

This goal can be achieved by addressing the following two scientific challenges.

Scientific challenge 1: Designing an effective quantum-to-classical distillation interface for robotics

The first challenge is to determine **what knowledge a quantum teacher should transfer** and **how to represent it** so that a classical student can learn it efficiently and still satisfy robotic constraints.

Unlike conventional teachers that output deterministic scores, a quantum teacher's outputs may vary due to finite measurement shots and hardware noise, and they depend strongly on circuit design choices (encoding, entanglement pattern, depth) [2]. This raises the central question of quantum-teacher distillation in robotics: which teacher signals best improve a student under realistic quantum constraints while preserving real-time deployability ?

A systematic comparison of these signals will clarify what information is most transferable and cost-effective.

Scientific challenge 2: Demonstrating data efficiency, robustness, and deployability in robotic settings

The second challenge is to validate whether the distilled student achieves measurable benefits compared to strong classical baselines, and whether those benefits persist when the quantum teacher is noisy or resource-limited.

In robotics, benefits must translate into operational gains such as lower end-to-end latency, higher reliability under changing conditions, and improved safety margins. We will target representative robotic learning problems where compact models matter:

- perception or state estimation from low-dimensional sensor features (e.g., fused proprioceptive descriptors and lightweight vision features);
- prediction of short-horizon dynamics, collision risk, or constraint violations;
- policy or value approximation components used inside decision-making or planning pipelines.

A prospective approach is to train quantum teachers (variational quantum classifiers and/or quantum-kernel teachers) on datasets aligned with near-term device constraints [3, 4], distill their outputs into compact classical students suitable for embedded deployment, and evaluate performance under conditions common in robotics: low-label regimes, sensor noise, and distribution shift. In addition to standard predictive metrics, we will evaluate calibration and uncertainty quality, since uncertainty estimates can support safer robot decision-making and monitoring.

A key outcome is a resource / performance tradeoff analysis: the student benefit as a function of quantum compute (shots, circuit evaluations, circuit depth), alongside system-oriented metrics such as inference latency, model size, and reliability under perturbations.

3 Work Plan

The work of this postdoc includes:

- Studying the state-of-the-art in knowledge distillation and quantum machine learning [5, 2]
- Implementing one or more quantum teacher models on quantum hardware and/or high-fidelity simulators [3, 4]
- Designing and benchmarking multiple quantum-to-classical distillation signals relevant to robotics applications
- Selecting robotics-relevant benchmarks and datasets and defining operational metrics (latency, energy, robustness, calibration)
- Evaluating distilled students in robotics-oriented settings such as low-label regimes, sensor noise, and distribution shift, reporting system constraints alongside accuracy
- Training and evaluating deployable classical students against strong classical baselines (including classical-teacher distillation controls)
- Producing a resource / performance analysis and publishing results in international conferences and journals

4 Skills

The candidate should preferably have a PhD-level background in quantum theory or machine learning, with strong experience in at least two of the following areas:

- Knowledge distillation / model compression
- Quantum machine learning (quantum kernels, variational quantum circuits)
- Experimental ML methodology (ablations, calibration, robustness, reproducibility)
- Robotics/Autonomy experience (perception, prediction, planning/control integration)

Proficiency in Python and modern ML tooling is expected, along with familiarity with a quantum SDK (e.g. Qiskit).

5 How To Apply

Interested candidates must send to Ezio Malis at `ezio.malis@inria.fr` the following documents:

- Motivation letter
- Letter of recommendation of the PhD thesis supervisor
- Curriculum vitæ including the list of the scientific publications

References

[1] B. Paden, M. Čáp, S. Z. Yong, D. Yershov, and E. Frazzoli, "A survey of motion planning and control techniques for self-driving urban vehicles," *IEEE Transactions on Intelligent Vehicles*, vol. 1, no. 1, pp. 33–55, 2016.

[2] M. Cerezo, G. Verdon, H.-Y. Huang, L. Cincio, and P. J. Coles, "Challenges and opportunities in quantum machine learning," *Nature Computational Science*, vol. 2, no. 9, pp. 567–576, 2022.

[3] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, "Supervised learning with quantum-enhanced feature spaces," *Nature*, vol. 567, no. 7747, pp. 209–212, 2019.

[4] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, "Parameterized quantum circuits as machine learning models," *Quantum Science and Technology*, vol. 4, no. 4, p. 043001, 2019.

[5] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," 2015.