ESANN 2020 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Online event, 2-4 October 2020, i6doc.com publ., ISBN 978-2-87587-074-2.
Available from http://www.i6doc.com/en/.

Machine learning framework for control in classical and quantum domains

Archismita Dalal, Eduardo J. Páez, Seyed Shakib Vedaie and Barry C. Sanders

Institute for Quantum Science and Technology,
University of Calgary, Alberta T2N 1N4, Canada

Abstract. Our aim is to construct a framework that relates learning and control for both classical and quantum domains. We claim to enhance the control toolkit and help un-confuse the interdisciplinary field of machine learning for control. Our framework highlights new research directions in areas connecting learning and control. The novelty of our work lies in unifying quantum and classical control and learning theories, aided by pictorial representations of current knowledge of these disparate areas. As an application of our proposed framework, we cast the quantum-control problem of adaptive phase estimation as a learning problem.

1 Introduction

The fields of classical control ($\mathscr{C}$C) [1] and quantum control ($\mathscr{Q}$C) [2] have been treated separately and have evolved independently. Some techniques from $\mathscr{C}$C theory have been successfully employed in $\mathscr{Q}$C, but lack of a formalised structure impedes progress of these techniques. Classical machine learning (ML) [3] and its emergent quantum counterpart [4] are prolific research fields with connections to the control field. Although these four research areas are well studied individually, we require a unified framework to un-confuse the community and enable researchers to see new tools. We develop a framework that relates learning and control for both classical and quantum domains, and assists in understanding how to use learning for control.

To generate a unified framework, we first review existing literature on control and learning, for both classical and quantum realms, emphasizing inconsistencies in definitions and explanations. We then elevate existing concepts on $\mathscr{C}$C to account for quantum aspects. Learning for control is built into our unified control scheme, with provisions for the various learning nuances. The aggregated literature and our framework are diagrammatically depicted for clarity. As an example of our framework, we demonstrate how to cast the $\mathscr{Q}$C problem of adaptive phase estimation (APE) into a supervised learning (SL) problem.

Our paper is organized as follows: In §2, we introduce the relevant background on control, ML and their connections. Here we briefly describe the $\mathscr{Q}$C problem that is presented as an example of our framework [§3] later in §4. Finally, §5 concludes the paper with future outlook and ideas.

ESANN 2020 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Online event, 2-4 October 2020, i6doc.com publ., ISBN 978-2-87587-074-2. Available from http://www.i6doc.com/en/.

2 Background

In this section we give a pertinent review of the fields of ML, control and their connections. In the last subsection we discuss the physical control problem of APE.

2.1 Machine learning

In 1997, Mitchell defined learning according to [3]

> An agent A is said to learn from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E.

This definition formally states the different components of a learning task, namely A, E, P and T, which are used to categorize a ML problem and subsequently apply the corresponding learning algorithm.

A ML task can be classified into three major categories based on the nature of the E it uses, namely SL, unsupervised learning and reinforcement learning. Another way to classify a ML task is based on the frequency of E of the agent. In an online learning scenario, an agent learns by generating a training data dynamically by interacting with the system in real-time; whereas in offline learning the agent starts with a static dataset.

Classical computational hardware cannot cope with the increasing volume of data being used in various ML applications in terms of their space and time complexities. Quantum computing, which has a fundamental advantage over classical computing in terms of processing speed, can be beneficial for such cases. Much research is done in the field of quantum-enhanced ML [4, 5], with a focus on classical-quantum hybrid learning algorithms to be implemented in near-time quantum devices [6, 7, 8]. On the other hand, ML helps with advancing quantum technologies, like $\mathscr{Q}$C [9, 10].

2.2 Control

The task of control is

> to use a controller C to steer specific controllable degrees of freedom of a plant P, consisting of the physical system with its input and output devices, such that its dynamics yields approximately the desired observations, specified by r.

In a typical feedback control scheme, the controller can follow a closed-loop strategy, whose actions on the plant are mediated by the control signal u, the information gained from the output y of the plant and a predetermined control policy ρ. This scheme is depicted in Fig. 1a. The control scheme can also be open-loop when the feedback from the plant is unavailable or unnecessary.

The idea of ML for control has been known since Fu's seminal concept of learning control for classical systems [11]. Learning enables a controller, who

ESANN 2020 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Online event, 2-4 October 2020, i6doc.com publ., ISBN 978-2-87587-074-2.
Available from http://www.i6doc.com/en/.

is neither omniscient nor possesses a feasible alternative, to execute the task successfully. Heuristic optimization algorithms, like genetic algorithms, enable searches for feasible policies [12]. On the other hand, ML can also help in modelling the plant's dynamics. In most practical scenarios, plant dynamics is complex and modelling including noise and fluctuations is infeasible. Learning the controller's policy is instead easier. Hybrid control schemes with learning for both plant and controller have also been studied [13].

Conventionally, the term $\mathscr{Q}$C is used when the dynamics of the plant is governed by quantum mechanics, and $\mathscr{C}$C refers to the case in which classical physics can accurately describe the dynamics of the plant [14, 15]. We can define open and closed-loop $\mathscr{Q}$C similar to their classical counterparts but with some subtleties arising due to the quantum measurement problem.

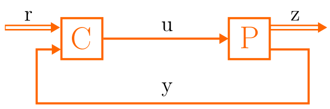

(a) Closed-loop $\mathscr{C}$C scheme

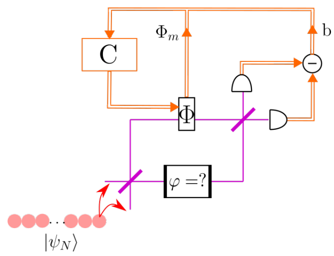

(b) APE scheme

Fig. 1

2.3 Adaptive quantum enhanced metrology

Adaptive quantum enhanced metrology (AQEM) is one strategy for performing quantum enhanced metrology, which aims to infer an unknown physical parameter using quantum resources, such that the scaling in uncertainty of that parameter surpasses the standard quantum limit (SQL) and approaches the Heisenberg limit [16]. In AQEM the control parameters are adaptively changed from one experiment to the next based on the measurement outcomes. In our discussion we consider the problem of optical APE, as the simplest case of AQEM, with applications to gravitational wave detection and atomic clocks.

Optical phase estimation in a Mach-Zehnder interferometer [Fig. 1b] involves an unknown phase shift φ, whose value has to be estimated, and a controllable phase shift Φ, whose value after the mth measurement is Φ_m. The interferometer has two input ports and two output ports. The input is an N-photon entangled state $|\psi_N\rangle$, which is sent to the interferometer one-photon-at-a-time. The output is denoted by $b_m \in \{0, 1\}$ according to the output port through which the mth photon exits the system. The controller C follows a policy ϱ to change Φ after each measurement such that it eventually approximates the unknown phase shift.

The $\mathscr{Q}$C problem in APE is to devise a feasible policy that delivers an imprecision in φ with scaling greater than SQL, based on the measurement outputs [9].

ESANN 2020 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Online event, 2-4 October 2020, i6doc.com publ., ISBN 978-2-87587-074-2.
Available from http://www.i6doc.com/en/.

The policy is a set of rules that determines how the controller adjusts the value of Φ_m in the mth round of measurement based on the previous measurement outcomes $\{b_1, b_2, \ldots, b_m\}$. For this work we are restricted to Markovian feedback with update rule

$$\Phi_m = \Phi_{m-1} - (-1)^{b_m} \Delta_m, \tag{1}$$

where ϱ is the ordered set $(\Delta_1, \Delta_2, \ldots, \Delta_N)$ and $\Delta_m \in [0, 2\pi)$. When all photons are exhausted, the estimate $\tilde{\varphi}$ is given by Φ_M, where $0 \leq M \leq N$ considering the case of photon loss. The imprecision $\Delta\tilde{\varphi}$ is quantified by the Holevo variance

$$V_N^{\varrho} := (S_N^{\varrho})^{-2} - 1, \tag{2}$$

where $S_N^{\varrho} := \left| \sum_{k=1}^{10N^2} \frac{\exp \mathrm{i}\left(\varphi^{(k)} - (\Phi_N^{\varrho})^{(k)}\right)}{10N^2} \right|$, and is minimized for the choice of a feasible policy.

3 Machine learning framework for control

In this section we first provide a graphical representation of the existing literature in the fields of control, learning and their connections. We then present our unified framework of ML for control and elaborate on its scope.

We begin by constructing a graphical representation for the background on learning and control, in the classical and quantum domains, to identify the existing and missing works in these fields [4, 17, 18, 19]. The graph is a directed square graph where the vertices represent the four fields, the horizontal edges define classical/quantum unification concepts, and vertical and diagonal edges define the idea of learning for control and vice-versa. We then identify authoritative sources in these fields that we treat as canons in our framework. In order to develop our unified concepts of control (learning), we use those canons but relax any restriction arising from classical physics. The square graph in Fig. 2a is obtained as a by-product of our proposed framework.

The diagrammatic representation of our proposed ML framework for control is provided in Fig. 2b. We introduce a teacher (T)/user (U), which is a classical agent (depicted in orange) governing the process of learning for control, and is different from the teacher in [11]. T determines the type of learning and instructs C, P and the learner L to act in the training or testing phase. We represent L, which executes the learning algorithm, in a purple box to account for the fact that it can be either classical or quantum. We use the same colour convention for the plant, controller and the various channels interconnecting them. The application of our framework is elaborated in the next section with a simple $\mathscr{Q}$C problem.

4 APE as a SL problem

As a demonstration of our proposed framework of ML for control, we describe in this section the control problem of APE as a SL problem and explain the stages of learning.

ESANN 2020 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Online event, 2-4 October 2020, i6doc.com publ., ISBN 978-2-87587-074-2.
Available from http://www.i6doc.com/en/.

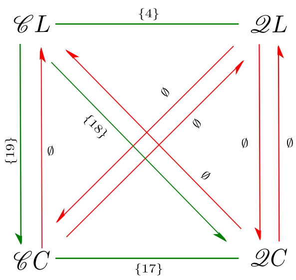

(a) Unifying framework: $\mathscr{C}$L and $\mathscr{Q}$L denotes classical machine learning and quantum machine learning respectively

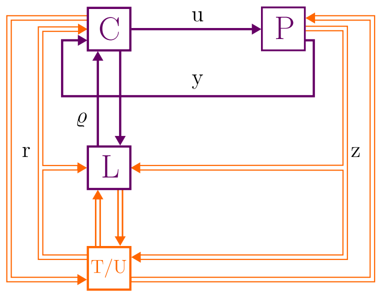

(b) Quantum/Classical learning control

Fig. 2

The algorithm for ML for APE starts with a decision by T whether to act in training or testing stage. In the training stage, T instructs both L and C to interact with the plant P to generate training data online and subsequently learn a feasible control policy ϱ. T also interacts with P to set its initial state. In the policy testing stage, T instructs just the controller to implement the learned policy on the plant to achieve the desired result.

We identify the Mach-Zehnder interferometric system, along with its input and output ports and detectors, of Fig. 1b as the plant P in our learning for control framework [Fig. 2b]. The input state $|\psi_N\rangle$ and the unknown phase φ of P is set by T. The controller sends signal u in the form of controllable phase shift Φ_m, which is calculated using the output b_m of P and a policy ϱ [Eq. (1)].

The learning agent L employs a SL algorithm to design a feasible control policy for estimating the unknown phase φ. To this end we map the different components of the $\mathscr{Q}$C problem to a SL problem. The training dataset $\{\boldsymbol{b}, \varphi\}$ is constructed online, where the feature vector $\boldsymbol{b}$ is a N-bit string corresponding to the sequential measurement outcomes. The label φ is the value of the unknown phase which is chosen randomly from $[0, 2\pi)$ during the training phase. The learning algorithm seeks a function $\Phi^{\varrho} : B \to \{\varphi\}$, where B is the space of $\boldsymbol{b}$ such that the Holevo variance [Eq. (2)] is minimized.

5 Conclusions

In this work we have reviewed state-of-the-art in learning for quantum and classical control. We have proposed a framework that unifies ML techniques for $\mathscr{Q}$C and $\mathscr{Q}$C. We show that APE fits into our proposed framework, demonstrating the viability of this new approach that might help the community to understand better the application of learning algorithms for control. During the preparation of this article we have identified the lack of works connecting quantum learning to $\mathscr{Q}$C and quantum learning to $\mathscr{C}$C.

ESANN 2020 proceedings, European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning. Online event, 2-4 October 2020, i6doc.com publ., ISBN 978-2-87587-074-2.
Available from http://www.i6doc.com/en/.

References

[1] Zhong-Sheng Hou and Zhuo Wang. From model-based control to data-driven control: Survey, classification and perspective. *Information Sciences*, 235:3–35, 2013.

[2] Howard M Wiseman and Gerard J Milburn. *Quantum measurement and control*. Cambridge university press, 2009.

[3] Thomas M. Mitchell. *Machine Learning*. McGraw-Hill, Inc., New York, NY, USA, 1 edition, 1997.

[4] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *Nature*, 549(7671):195, 2017.

[5] Vedran Dunjko and Hans J Briegel. Machine learning & artificial intelligence in the quantum domain: a review of recent progress. *Reports on Progress in Physics*, 81(7):074001, 2018.

[6] JS Otterbach, R Manenti, N Alidoust, A Bestwick, M Block, B Bloom, S Caldwell, N Didier, E Schuyler Fried, S Hong, et al. Unsupervised machine learning on a hybrid quantum computer. *arXiv preprint arXiv:1712.05771*, 2017.

[7] Anna Levit, Daniel Crawford, Navid Ghadermarzy, Jaspreet S Oberoi, Ehsan Zahedinejad, and Pooya Ronagh. Free energy-based reinforcement learning using a quantum processor. *arXiv preprint arXiv:1706.00074*, 2017.

[8] Moslem Noori, Seyed Shakib Vedaie, Inderpreet Singh, Daniel Crawford, Jaspreet S Oberoi, Barry C Sanders, and Ehsan Zahedinejad. Adiabatic quantum kitchen sinks for learning kernels using randomized features. *arXiv preprint arXiv:1909.08131*, 2019.

[9] Pantita Palittapongarnpim, Peter Wittek, Ehsan Zahedinejad, Shakib Vedaie, and Barry C Sanders. Learning in quantum control: High-dimensional global optimization for noisy quantum dynamics. *Neurocomputing*, 268:116–126, 2017.

[10] Ehsan Zahedinejad, Joydip Ghosh, and Barry C Sanders. Designing high-fidelity single-shot three-qubit gates: a machine-learning approach. *Physical Review Applied*, 6(5):054005, 2016.

[11] King-Sun Fu. Learning control systems–review and outlook. *IEEE transactions on Automatic Control*, 15(2):210–221, 1970.

[12] Dimitris C Dracopoulos and Simon Kent. Genetic programming for prediction and control. *Neural Computing & Applications*, 6(4):214–228, 1997.

[13] Thomas Duriez, Steven L Brunton, and Bernd R Noack. *Machine Learning Control-Taming Nonlinear Dynamics and Turbulence*. Springer, 2017.

[14] Andrew C Doherty, Salman Habib, Kurt Jacobs, Hideo Mabuchi, and Sze M Tan. Quantum feedback control and classical control theory. *Physical Review A*, 62(1):012105, 2000.

[15] Daoyi Dong and Ian R Petersen. Quantum control theory and applications: a survey. *IET Control Theory & Applications*, 4(12):2651–2671, 2010.

[16] Géza Tóth and Iagoba Apellaniz. Quantum metrology from a quantum information science perspective. *Journal of Physics A: Mathematical and Theoretical*, 47(42):424006, 2014.

[17] Seyed Shakib Vedaie, Pantita Palittapongarnpim, and Barry C Sanders. Reinforcement learning for quantum metrology via quantum control. In *2018 IEEE Photonics Society Summer Topical Meeting Series (SUM)*, pages 163–164. IEEE, 2018.

[18] Jing Zhang, Yu-xi Liu, Re-Bing Wu, Kurt Jacobs, and Franco Nori. Quantum feedback: theory, experiments, and applications. *Physics Reports*, 679:1–60, 2017.

[19] Karl Johan Åström. Theory and applications of adaptive control—a survey. *Automatica*, 19(5):471–486, 1983.