International Journal of Quantum Information
(2024) 2450002 (6 pages)

DOI: 10.1142/S0219749924500023

Application of machine learning to experimental design in quantum mechanics

Federico Belliardo*
NEST, Scuola Normale Superiore, I-56126 Pisa, Italy
federico.belliardo@sns.it

Fabio Zoratti
Scuola Normale Superiore, I-56126 Pisa, Italy

Vittorio Giovannetti
NEST, Scuola Normale Superiore and Istituto Nanoscienze-CNR
I-56126 Pisa, Italy

Received 26 December 2023
Accepted 9 January 2024
Published

The recent advances in Machine Learning hold great promises for the field of quantum sensing and metrology. With the help of reinforcement learning, we can tame the complexity of quantum systems and solve the problem of optimal experimental design. Reinforcement learning is a powerful model-free technique that allows an agent, which is typically a neural network, to learn the best strategy to reach a certain goal in a completely a priori unknown environment. However, in general, we know something about the quantum system the agent is interacting with, at least that it follows the rules of quantum mechanics. In quantum metrology, we typically have a model for the system and only some parameters of the evolution or of the initial state are unknown. We present here a general Machine Learning technique that can optimize the precision of quantum sensors, and in doing so it exploits the knowledge we have on the system. We have developed a Python package to automate a broad class of optimizations that can be found in the tasks of quantum parameter estimation, quantum metrology and quantum hypothesis testing. What the agent is learning here is an optimal adaptive strategy, that, on the basis of the previous outcomes, decides the next measurements to perform. It works both for Bayesian estimation and for frequentist estimation. The user is required to implement the physics of the system to be studied and state which parameters in the experiment are controllable and which are unknowns. The functions of the library allow then to easily train a neural network agent for optimizing the precision of the sensor, by simulating the experiment. We have explored some applications of this technique to magnetometry on NV centers (both DC and AC), to state discrimination in quantum optics and to phase estimation. So far, we were able to certify better results than the current state-of-the-art controls for many examples. The Machine Learning

*Corresponding author.

technique developed here can be applied in all those scenarios where the quantum system is well characterized and relatively simple and small. In these cases, we can squeeze every last bit of information from a quantum sensor by controlling it with a neural network appropriately trained.

Keywords: Machine learning; quantum metrology; quantum sensing; nv center; reinforcement learning.

1. Introduction

In recent times, there has been a growing focus on the intersection of Machine Learning and quantum information. The collaboration between these two technological realms holds promise for mutual benefits. Quantum technologies, particularly quantum computers, possess the capability to tackle conventional challenges in Machine Learning, such as classification and pattern recognition, whether handling classical or quantum data.[1,2] Conversely, conventional Machine Learning can enhance tasks in quantum information, such as quantum control with feedback[3] and error correction.[4] Our research falls into the latter category. Specifically, we employ model-aware reinforcement learning to discover optimized adaptive and nonadaptive control strategies for tasks in quantum metrology, estimation and hypothesis testing. Through this approach, we investigate how Machine Learning has the potential to improve traditional methods in quantum physics and contribute to the advancement of new quantum information processing technologies. In quantum metrology experiment we are interest in the estimation of some unknown parameters and the goodness of the experiment and the data processing can be gauged by an error figure of merit, e.g. the mean square error relative to the true values of the unknown parameters. After specifying a set of adjustable variables within an experiment, an agent employs reinforcement learning to effectively manipulate them and minimize the error metric. The agent can take the form of a compact neural network, a decision tree, or a straightforward list of trainable controls applied sequentially. The whole controlled estimation has been abstracted from the specific sensor and physical platform and encapsulated into the qsensoropt library,[5] which will soon be accessible on PyPI. Consequently, this library serves as a versatile tool for optimizing a diverse range of quantum sensors. Our framework was tested across various examples using the nitrogen-vacancy (NV) center platform,[6,7] encompassing DC[8] and AC magnetometry, decoherence estimation,[9] and hyperfine coupling characterization.[10] We report as example the estimation of the hyperfine coupling of the electron spin in a NV center with a ^{13}C in the carbon lattice, see Fig. 1. Additionally, within the realm of photonic circuits, we explored multiphase hypothesis testing, a recent extension of the Dolinar receiver,[11] and its adaptation to the discrimination of three states, along with coherent states classification. In the frequentist estimation domain, our investigation focused on detuning frequency sensing in a driven optical cavity.[12] Comprehensive details regarding the library's implementation and these diverse examples can be found in Ref. 13 in the library's online documentation. Our results demonstrate that reinforcement learning surpasses traditional

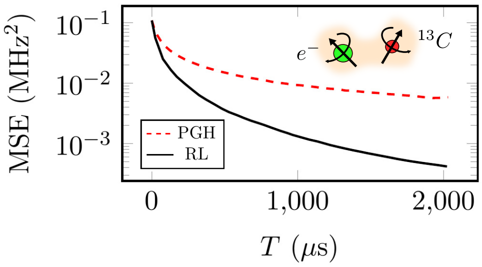

Fig. 1. Mean square error (MSE) for the estimation of the hyperfine dipolar coupling between the electron spin and a ^{13}C nucleus, obtained controlling the MW pulses with the Particle Guess Heuristic (PGH) used in Ref. 10, compared to the performances of model-aware RL. The precision is reported as a function of the total free precession time T of the electronic spin in an external magnetic field.

control strategies across multiple scenarios. This research lays the groundwork for accelerating the quest for optimal controls in quantum sensors, potentially accelerating their widespread industrial application.

2. The Measurement Loop

In quantum metrology, estimation and hypothesis testing, we analyze a physical system, called quantum probe, governed by a well-known quantum dynamic, which the experimenter can continuously modify by selecting values for a fixed set of controls (e.g. a tunable phase in an interferometer). The experimenter is interested in estimating certain unknown parameters, denoted by $\boldsymbol{\theta}$, associated with the environment and encoded in the probe through its interaction, or associated to the initial state of the probe itself. These parameters are estimated through measurements on the probe. The estimation process employs a particle filter[14–16] (PF) to process the outcomes of the measurements, which uses the Bayes' rule to update the Bayesian posterior probability distribution on the parameters $\boldsymbol{\theta}$ after each measurement. The PF employs a set of discrete samples, or particles, with a weight associated to each of them, to represent the posterior probability distribution. Controls are determined based on information within the PF, such as the mean and variance of the distribution, which are the input to the agent that produces the controls as output. The next measurement is then performed and the process is repeated to form the *measurement loop*, of which a single iteration is represented in Fig. 2. The knowledge of the parameters $\boldsymbol{\theta}$ in the PF obtained through measurements is leveraged to guide the evolution and measurements on the probe, optimizing the overall performance of the estimation task.

3. The Sensors Model

While the Bayesian filtering, the probe's control and the agent's training have been implemented in the library, the user must implement a differentiable model of the

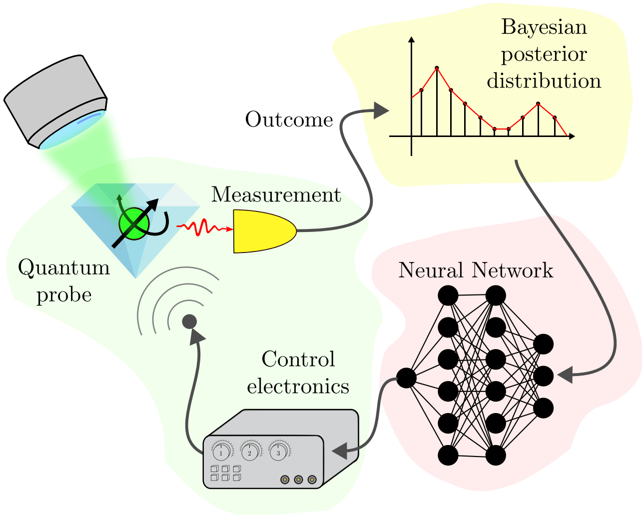

Fig. 2. This general scheme illustrates the information flow within the measurement loop. The environment we aim to study interacts with the quantum probe and encodes it with the unknown variables $\boldsymbol{\theta}$. This probe is then measured using a tunable instrument. The outcome of this measurement provides us with information about the probe's state and in turn about the environment's variables. This information is used in the particle filter to update the posterior Bayesian distribution on $\boldsymbol{\theta}$. Some summary information derived from the particle filter is then input into an agent that decides the new control parameters for the measurement in the next iteration of the loop. This control is then realized through the electronics of the experiment. In this picture, the agent is a neural network.

sensor of their interests using TensorFlow. This model should simulate the stochastic extraction of measurement outcomes and evaluate the probability of observing a specific outcome in a measurement. The reinforcement learning approach enables us to address a diverse range of tasks using a unified tool.

4. The Precision Resource-Paradigm

Every iteration of the measurement loop consumes some amount of a specific "resource", which is costly in the context of the estimation and must be defined by the user. Once these resources are depleted, the measurement loop concludes. Examples of resources are the total estimation time, the number of measurement performed or the intensity of a signal.

5. Training Loss

Once the estimation ends and the measurement loop is terminated an estimator $\widehat{\boldsymbol{\theta}}$ for the parameters $\boldsymbol{\theta}$ is computed fro the PF. From this estimator, the user-defined precision metric for the sensor is evaluated. This might be, for instance, the mean square error in a parameter estimation task or the error probability in an hypothesis testing scenario. This precision metric is the loss to be minimized.

6. Training of the Agent

The training of the agent for each experiment is facilitated by the functions of our library. For the majority of the applications, we opted for neural networks due to their proven suitability to approximate generic smooth functions.[17] Through automatic differentiation over all iterations of the loop, a gradient descent procedure is employed to train the agent to minimize the loss. The gradient flows through the stochastic measurement outcome extraction, the update of the probe state and the Bayes' rule applied to the particle filter ensemble. By default, the gradient doesn't propagate through the agent's input. Allowing the gradient to traverse the physical model implementation of the sensor categorizes this training as a form of model-aware policy gradient reinforcement learning. Since the loss is a stochastic variable, as it depends on the simulated measurement outcomes, special precautions are necessary to compute an unbiased estimator for its gradient.[3]

7. Conclusions

In summary, our research underscores the advantages of integrating Machine Learning with contemporary quantum technologies. We have introduced a framework, complemented by a versatile library, designed to address a broad range of challenges in quantum parameter estimation and metrology. This library provides a flexible interface, allowing researchers to easily configure and optimize diverse parameter estimation tasks based on quantum systems. With the potential to expedite the development of practical applications in quantum parameter estimation and metrology, our library opens avenues for precise estimation of physical parameters that could transform various sectors, including biology, fundamental physics and quantum communication. By offering a user-friendly tool, we aim to facilitate progress in these domains, facilitating the transition of quantum-based metrology from proof-of-principle experiments to industrial applications.

Acknowledgments

We gratefully acknowledge the computational resources of the Center for High Performance Computing (CHPC) at SNS. We acknowledge financial support by MUR (Ministero dell'Istruzione, dell'Universitá e della Ricerca) through the following projects: PNRR MUR project PE0000023-NQSTI, PRIN 2017 Taming complexity via Quantum Strategies: a Hybrid Integrated Photonic approach (QU-SHIP) Id. 2017SRN-BRK.

ORCID

Federico Belliardo https://orcid.org/0000-0002-1466-396X
Fabio Zoratti https://orcid.org/0000-0003-2390-0602
Vittorio Giovannetti https://orcid.org/0000-0002-7636-9002

References

1. M. Broughton *et al.*, arXiv:2003.02989.
2. V. Bergholm *et al.*, arXiv:1811.04968.
3. R. Porotti, V. Peano and F. Marquardt, *PRX Quantum* **4** (2023) 030305.
4. T. Fsel, P. Tighineanu, T. Weiss and F. Marquardt, *Phys. Rev. X* **8** (2018) 031084.
5. F. Belliardo and F. Zoratti, Qsensoropt GitLab repository, https://gitlab.com/federico.belliardo/qsensoropt (2024).
6. M. Chen *et al.*, *Natl. Sci. Rev.* **5** (2018) 346.
7. P. Rembold *et al.*, *AVS Quantum Sci.* **2** (2020) 024701.
8. L. J. Fiderer, J. Schuff and D. Braun, *PRX Quantum* **2** (2021) 020303.
9. M. J. Arshad *et al.*, arXiv:2210.06103.
10. T. Joas, *npj Quantum Inf.* **7** (2021) 56.
11. F. Zoratti, N. D. Pozza, M. Fanizza and V. Giovannetti, *Phys. Rev. A* **104** (2021) 042606.
12. A. Fallani, M. A. C. Rossi, D. Tamascelli and M. G. Genoni, *PRX Quantum* **3** (2022) 020310.
13. F. Belliardo, F. Zoratti, F. Marquardt and V. Giovannetti, arXiv:2312.16985.
14. P. Del Moral, *C. R. Acad. Sci. I* **325** (1997) 653.
15. M. Arulampalam, S. Maskell, N. Gordon and T. Clapp, *IEEE Trans. Signal Process.* **50** (2002) 174.
16. J. S. Liu and R. Chen, *J. Am. Stat. Assoc.* **93** (1998) 1032.
17. T. De Ryck, S. Lanthaler and S. Mishra, *Neural Netw.* **143** (2021) 732.