Workshop Summary: Quantum Machine Learning

Volker Tresp
Ludwig Maximilian University
Siemens AG, Technology
Munich, Germany
volker.tresp@lmu.de

Steffen Udluft
Siemens AG, Technology
Munich, Germany
steffen.udluft@siemens.com

Daniel Hein
Siemens AG, Technology
Munich, Germany
hein.daniel@siemens.com

Werner Hauptmann
Siemens AG, Technology
Munich, Germany
werner.hauptmann@siemens.com

Martin Leib
IQM
Munich, Germany
martin.leib@meetiqm.com

Christopher Mutschler
Fraunhofer IIS
Nuremberg, Germany
christopher.mutschler@iis.fhg.de

Daniel D. Scherer
Fraunhofer IIS
Nuremberg, Germany
daniel.scherer2@iis.fhg.de

Wolfgang Mauerer
Technical University of
Applied Sciences Regensburg
Siemens AG, Technology
Regensburg/Munich, Germany
wolfgang.mauerer@othr.de

I. Workshop Rationale

Quantum computing (QC) has made significant progress in recent years, and scientists are exploring its applications across various fields, including quantum machine learning (QML).

The workshop 'Quantum Machine Learning: From Foundations to Applications', part of the IEEE International Conference on Quantum Computing and Engineering 2023, aims to bring together researchers and industry practitioners from different disciplines to discuss the challenges and applications of QML.

Quantum computing (QC) has seen remarkable progress after first commercial hardware prototypes have become available in recent years, following decades of foundational research. Scientists across all fields are seeking concrete applications for the technology. Fault-tolerant and large-scale quantum computers are still unavailable, yet many opportunities beckon for machine learning: Given reasonable complexity-theoretic assumptions, quantum algorithms are expected to be capable of outperforming known approaches to problems even on noisy, intermediate-scale quantum (NISQ) computers, and exhibit advantages over classical systems.

For the domains of machine learning and artificial intelligence, QC is a relatively new technology, and the exact impact on the field is still unclear. The goal of our workshop is to unite academic researchers and industry practitioners from multiple disciplines (*e.g.*, AI, ML, software/systems engineering, physics, etc.) to discuss challenges and applications of quantum machine learning (including machine learning to improve quantum computing itself) that can advance the state of the art. Apart from enabling interaction between AI/ML and the quantum computing research communities, we intend to closely incorporate industrial users to identify application potentials, as well as to explore co-design ideas that enable special-purpose, hybrid quantum-classical appliances to be designed for problems of topical importance.

Many machine learning techniques (see, *e.g.*, Refs. [5], [1], [24], [2], [16]) are backed by quantum counterparts. Seminal results show that, for instance, quantum support vector machines [17] can achieve exponential speedups over classical approaches. Likewise, quantum ML could operate on fewer data points than classical methods [3], while quantum classifiers [9], which have even been implemented on small-scale NISQ machines, can derive hard to (classically) estimate kernels. Topological data analysis [11], PCA [12], or relational learning on knowledge graphs [13] further augment the list of quantum-accelerated ML tasks. Yet, the conditions of their advantages need careful consideration of subtle issues (see, *e.g.*, Refs. [22], [19], [7], [14]) that do not appear in classical approaches. Concrete *practical* applications of QML are still unknown. Given the interdisciplinary and cross-domain nature of the field, there is still a lack of opportunities to publish early results that can benefit from interactive, engaged discussion between experts at a workshop. QCE presents an ideal opportunity to create an interdisciplinary dialogue between experts of diverse fields. It enables to present the young but thriving field of QML to a huge community.

Machine learning, as one of the most active research fields at present, encompasses a huge community. Therefore, it is a pertinent question if the topic of quantum machine learning should better be placed at a quantum-focused or an ML-focused conference. Lighthouse conferences for machine learning (like ICML or NeurIPS) have grown considerably in the last years. The practical maturity of quantum machine learning techniques is substantially less advanced than their canonical objects of study, and especially cannot compete with current solutions in terms of scalability and performance, given their decades worth of engineering head-start. Therefore, we believe that quantum machine learning should be better addressed at a venue with quantum focus, where the topic occupies less of a niche.

Our long-term vision is to anchor the topic in the quantum community. This includes raising awareness for the specific challenges and opportunities associated with QML, and to allow for an active transfer of knowledge between the two involved communities. We also aim at improving synergies between industrial practitioners and researchers: While the former can

provide apt use-cases for QML, and gauge the aptitude of principal approaches in realistic scenarios, the latter can learn what directions of research are deemed promising for achieving eventual quantum superiority over classical approaches.

II. DESIRED CONTRIBUTIONS AND OUTCOMES

A. Contribution Format

The workshop invited submissions of papers in the following categories:

Research papers propose new approaches, theories, or techniques related to QML, including new algorithms. They should make substantial theoretical and empirical contributions to the field of research.

System papers describe new systems and accelerator approaches for QML, in particular through hardware-algorithm co-design, and with a focus on interaction of classical and quantum technology.

Experiments and analysis papers focus on the experimental evaluation of existing approaches, including methods and algorithms for QML, and bring new insights through the analysis of experiments and simulations. Results can, for example, show benefits of established approaches in a quantum setting or establish research needs by demonstrating unexpected behaviour.

Application papers report practical experience on applications of quantum machine learning based on sound scientific evaluation. This includes industrial experience reports.

B. Outcomes

The desired outcomes of our workshop are multifold:

- Publish, present and discuss novel research contributions covering all aspects of quantum machine learning.
- Allow participants to form interest groups around topics that may lead to cooperative work or grant applications.
- Improve the understanding of means to compare performance and benefits of QML solutions in a holistic way to established machine learning approaches, especially considering realistic industrial or practical constraints.
- Identify research opportunities and application domains for QML.
- Form industrial-academic relationships that can be the basis for future consortia.
- Exchange of ideas between machine learning and quantum computing communities.

III. TARGET AUDIENCE

The target audience for our workshop is everyone with an interest in learning about possibilities and challenges for quantum computing in machine learning. We will make efforts to attract contributions and participants from both, industry and academia. We explicitly encouraged submissions on the boundaries between software and hardware, and on co-design efforts between physical implementation and problem domain(s), to bring together experience from computer science, engineering, and physics.

IV. REVIEW PROCESS AND RESULT

A. Approach

We strived to implement a triple-anonymous peer review process (identities of author(s), reviewers and editor(s) are hidden from each other, and identities of authors and possibly editors are only revealed post acceptance) to the best of our abilities, given some challenges with the mandatory easychair setup of the umbrella conference, which mandated that some of the organisers (who were not part of the review process) needed to learn identities of submitting authors. All manuscripts enjoyed at least three reviews, evaluating by (a) relevance to the workshop, (b) novelty, practical impact and scientific depth, (c) technical soundness, (d) appropriateness and adequacy of literature coverage and result discussion, and (e) presentation, including overall organisation and readability.

The organisers initially estimated to accept about nine papers, in good agreement with eight eventually contributed papers.

B. Accepted Papers

From the eight accepted papers, half concern experiments and their analysis, while the other half present new ideas and approaches. With four papers originating from academic authors, two enjoying a mixed academic/industrial background, and two coming from industrial contributors, the target of joining academia and industry seems to have been achieved.

We accepted submissions from three continents (Europe, India, and North America), which is in line with the desired international character of the workshop (four papers originating from Germany might testify to the fact that the organisers have collaborated in a consortium funded by the German National Ministry for Research for over two years, thus inducing a slight bias on the outlets where the call for papers resonated with prospective authors).

Notably, all except one paper do *not* rely on experimentation on real quantum hardware, but perform numerical simulations on classical systems. We can, broadly, identify three main directions that contributions address:

Applications of QML. Both papers, 'A Quantum-annealing-based Transfer Learning Approach for Large-Scale Image Classification Using Quantum Boltzmann Machines' [18] and 'Applying QNLP to sentiment analysis in finance' [20] hypothesise improvements over existing classical baselines for the respective approaches based on numerical simulations, but cannot rigorously establish such advantages, owing to difficulties in extending results obtained on tractable smaller samples to general cases.

Quantum Reinforcement Learning. 'Quantum Natural Policy Gradients: Towards Sample-Efficient Reinforcement Learning' [15] introduces a quantum natural policy gradient algorithm that extends previously established quantum policy gradient approaches by incorporating second-order terms, which can in turn improve sampling efficiency. 'Differentiable Quantum Architecture Search for Quantum Reinforcement learning' [21] introduces automatic search techniques

to establish good quantum circuits for quantum deep-Q learning, and applies the method to two seminal QRL benchmark problems (frozen lake and cartpole). 'qgym: A Gym for Training and Benchmarking RL-Based Quantum Compilation' [23] uses (classical) reinforcement learning to optimise circuit compilation, which could in turn be applied to many quantum machine learning applications. 'Quantum deep Q learning with distributed prioritized experience replay' [4] extends QRL to asynchronous, multi-core capable distributed training, potentially involving multiple QPUs. Numerical simulations demonstrate performance improvements over previous quantum deep-Q network approaches (whose speedup potential over classical approaches is debated).

General Quantum Machine Learning. 'Towards an End-To-End Approach for Quantum Principal Component Analysis' [6] discusses a complete implementation of a seminal algorithm without simplifying assumptions or shortcuts whose objective—computing the principal components of a matrix—can be used in several machine learning algorithms. 'Quantum Machine Learning with Quantum Topological Data Analysis' [10] uses efficient quantum methods to compute Betti numbers and persistent Betti numbers, which have important known uses in classical ML, yet are hard to determine.

C. Evaluation

We believe that the accepted papers paint an accurate picture of the state of the field, which is highly dynamic and diverse in topics, as well as of the currently existing research challenges: Unconditional speedups for quantum machine learning over established classical approaches are not only extremely hard to obtain due to the lack of theoretical understanding of the former, but also very much owing to the wide mixture of heuristic approaches in the latter. The peculiarities of existing machines, especially their small scale, high level of noise and imperfections, among other limitations [8], make it hard to extrapolate larger-scale behaviour, or to consider the end-to-end properties of complete systems. Likewise, the strong dependence on physical foundations is unaccustomed to computer science. Compared to the extremely large number of active researchers in classical machine learning, the QML community is tiny, which is likely influenced by the considerable barrier that mastering two large domains—quantum computing and machine learning—entails. Nonetheless, we are confident that our workshop can serve to strengthen existing connections between research groups collaborating on the matter, initiate new collaborations, and draw the attention of quantum computing researchers to the many open challenges of quantum machine learning.

REFERENCES

[1] Z. Abohashima, M. Elhosen, E. H. Houssein, and W. M. Mohamed. Classification with quantum machine learning: A survey. *arXiv*, arXiv:2006.12270, 2020.

[2] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd. Quantum machine learning. *Nature*, 549(7671), 2017.

[3] M. C. Caro, H.-Y. Huang, M. Cerezo, K. Sharma, A. Sornborger, L. Cincio, and P. J. Coles. Generalization in quantum machine learning from few training data. *Nat. Commun.*, 13(1), 2022.

[4] S. Y.-C. Chen. Quantum deep Q learning with distributed prioritized experience replay. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 2023.

[5] C. Ciliberto, M. Herbster, A. D. Ialongo, M. Pontil, A. Rocchetto, S. Severini, and L. Wossnig. Quantum machine learning: a classical perspective. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 474(2209):20170551, 2018.

[6] E. Dri, A. Aita, T. Fioravanti, G. Franco, E. Giusto, G. Ranieri, B. Montrucchio, and D. Corbelletto. Towards an end-to-end approach for quantum principal component analysis. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 2023.

[7] M. Franz, L. Wolf, M. Periyasamy, C. Ufrecht, D. Scherer, A. Plinge, C. Mutschler, and W. Mauerer. Uncovering instabilities in variational-quantum deep Q-networks. *Journal of The Franklin Institute*, 2022.

[8] F. Greiwe, T. Krüger, and W. Mauerer. Effects of imperfections on quantum algorithms: A software engineering perspective. In *Proceedings of the IEEE International Conference on Quantum Software*, 2023.

[9] 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*, 567(7747):209–212, 2019.

[10] A. Khandelwal and M. G. Chandra. Quantum machine learning with quantum topological data analysis. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 2023.

[11] S. Lloyd, S. Garnerone, and P. Zanardi. Quantum algorithms for topological and geometric analysis of data. *Nature Communications*, 7(1):10138, 2016.

[12] S. Lloyd, M. Mohseni, and P. Rebentrost. Quantum principal component analysis. *Nature Physics*, 10(9):631–633, 2014.

[13] Y. Ma and V. Tresp. Quantum machine learning algorithm for knowledge graphs. *ACM Transactions on Quantum Computing*, 2(3), 2021.

[14] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1):4812, 2018.

[15] N. Meyer, D. D. Scherer, A. Plinge, C. Mutschler, and M. J. Hartmann. Quantum natural policy gradients: Towards sample-efficient reinforcement learning. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 2023.

[16] N. Meyer, C. Ufrecht, M. Periyasamy, D. D. Scherer, A. Plinge, and C. Mutschler. A survey on quantum reinforcement learning. 2022.

[17] P. Rebentrost, M. Mohseni, and S. Lloyd. Quantum support vector machine for big data classification. *Phys. Rev. Lett.*, 113:130503, 2014.

[18] D. J. Schuman, L. Sünkel, P. Altmann, J. Stein, C. Roch, T. Gabor, and C. Linnhoff-Popien. A quantum-annealing-based transfer learning approach for large-scale image classification using quantum Boltzmann machines. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 2023.

[19] A. Skolik, J. R. McClean, M. Mohseni, P. van der Smagt, and M. Leib. Layerwise learning for quantum neural networks. *Quantum Machine Intelligence*, 3(1):5, 2021.

[20] J. Stein, I. Christ, N. Kraus, M. B. Mansky, R. Müller, and C. Linnhof-Popien. Applying QNLP to sentiment analysis in finance. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 2023.

[21] Y. Sun, Y. Ma, and V. Tresp. Differentiable quantum architecture search for quantum reinforcement learning. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 2023.

[22] E. Tang. Quantum principal component analysis only achieves an exponential speedup because of its state preparation assumptions. *Phys. Rev. Lett.*, 127:060503, 2021.

[23] S. van der Linde, W. de Kok, T. Bontekoe, and S. Feld. qgym: A Gym for training and benchmarking RL-based quantum compilation. In *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 2023.

[24] R. Zhao and S. Wang. A review of quantum neural networks: Methods, models, dilemma. *arXiv*, 2109.01840, 2021.