Check for updates

Quantum Machine Learning: Foundation, New Techniques, and Opportunities for Database Research

Tobias Winker
Sven Groppe
University of Lübeck
Lübeck, Germany

Valter Uotila
Zhengtong Yan
Jiaheng Lu
University of Helsinki
Helsinki, Finland

Maja Franz
Wolfgang Mauerer
Technical University of Applied Science Regensburg
Regensburg, Germany

ABSTRACT

In the last few years, the field of quantum computing has experienced remarkable progress. The prototypes of quantum computers already exist and have been made available to users through cloud services (e.g., IBM Q experience, Google quantum AI, or Xanadu quantum cloud). While fault-tolerant and large-scale quantum computers are not available yet (and may not be for a long time, if ever), the potential of this new technology is undeniable. Quantum algorithms have the proven ability to either outperform classical approaches for several tasks, or are impossible to be efficiently simulated by classical means under reasonable complexity-theoretic assumptions. Even imperfect current-day technology is speculated to exhibit computational advantages over classical systems. Recent research is using quantum computers to solve machine learning tasks. Meanwhile, the database community has already successfully applied various machine learning algorithms for data management tasks, so combining the fields seems to be a promising endeavour. However, quantum machine learning is a new research field for most database researchers. In this tutorial, we provide a fundamental introduction to quantum computing and quantum machine learning and show the potential benefits and applications for database research. In addition, we demonstrate how to apply quantum machine learning to the join order optimization problem in databases.

CCS CONCEPTS

• **Computer systems organization → Quantum computing**; • **Computing methodologies → Machine learning**; • **Information systems → Data management systems**.

KEYWORDS

Quantum machine learning, quantum computing, databases

ACM Reference Format:
Tobias Winker, Sven Groppe, Valter Uotila, Zhengtong Yan, Jiaheng Lu, Maja Franz, and Wolfgang Mauerer. 2023. Quantum Machine Learning: Foundation, New Techniques, and Opportunities for Database Research. In *Companion of the 2023 International Conference on Management of Data (SIGMOD-Companion '23), June 18–23, 2023, Seattle, WA, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3555041.3589404

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGMOD-Companion '23, June 18–23, 2023, Seattle, WA, USA
© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9507-6/23/06…
https://doi.org/10.1145/3555041.3589404

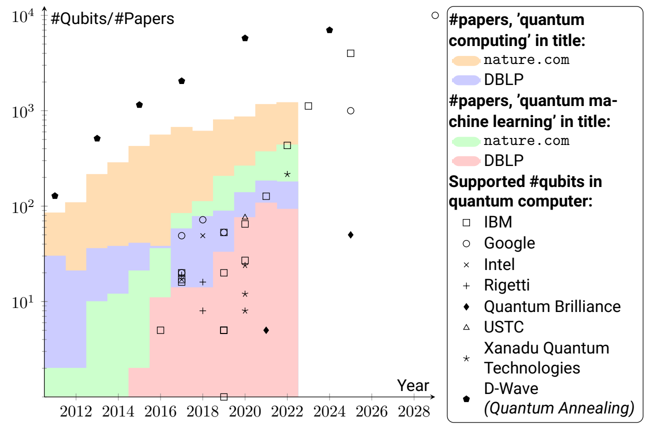

Figure 1: Timeline of quantum computing and quantum machine learning papers, and quantum computers (including roadmaps). Figure is extended from [24].

1 INTRODUCTION

Considering the timeline of available and future quantum computers in relation to the number of supported qubits in Figure 1, there seems to be an exponential growth trend in the number of supported qubits. The roadmap of major players contains quantum computers (QC) allowing to scale in 2023 (IBM), supporting 4000 qubits in 2025 (IBM) and 10000 qubits in 2029 (Google). Although the number of qubits is known to be a problematic measure for general QC capabilities (and other metrics such as quantum volume [13] have been proposed), the prestigious race for the most qubits is a driver of the current hype in quantum technologies promising numerous quantum applications in practice within this decade.

A quite obvious correlation exists between the availability of quantum computers supporting more qubits and the publication performance of researchers in the areas of quantum computing and quantum machine learning (see Figure 1). There seem to be differences in the absolute numbers of published papers in the addressed areas for different scientific communities: In 2022, there have been 6.8 times more papers published on nature.com (aiming to publish journal articles in the areas of natural sciences) containing 'quantum computing' and 4.7 times more papers containing 'quantum machine learning' in the title than are included in the dblp computer science bibliography (providing open bibliographic information on major computer science journals and proceedings). When searching for 'machine learning' in the title of papers published, it seems to be the other way round: In 2022, DBLP contains 2.3 times more papers than nature.com. Looking at no specific range of dates, then DBLP contains even 4.6 times more papers. With this tutorial, we want to encourage computer scientists (especially those with an interest in data management) to explore the possibilities of quantum computing for their research area.

Many research contributions (as summarized in surveys [2, 12, 69]) propose and discuss, for nearly every major approach of machine learning, a corresponding quantum counterpart. There has been a focus on investigating the benefits of applying quantum machine learning over classical machine learning. For example, applying quantum support vector machines [50] may achieve an exponential speedup in comparison to their classical variants. It has been shown that quantum machine learning methods may have the desirable advantage to learn on fewer data points than classical methods [10]. The properties of data sets with a potential quantum advantage in learning tasks have been identified in [30].

Recent contributions solve database problems including join order optimization [53, 55, 60, 63] and transaction scheduling [7, 8, 23] on quantum computers, but none of them applies quantum machine learning so far.

Tutorial overview: In this tutorial, after starting with the motivation describing the hype of quantum computing and quantum machine learning, we will provide a dive into the basics of quantum computing (see Section 2). Afterward, we will introduce quantum machine learning by comparing it with classical machine learning (see Section 3). In a hands-on tutorial, we will show how to apply quantum computing with a special focus on quantum machine learning for database tasks like join order optimization (see Section 4). Finally, we will discuss quantum machine learning opportunities and future work for database research (see Section 5).

Related tutorials: Previous tutorials in SIGMOD [32, 34, 49, 62] have already addressed the benefits of classical machine learning to data management tasks like query optimization, learned indices, workload prediction, cardinality and cost estimations, natural language interfaces to data, automating exploratory data analysis and data cleaning. In principle, all these studied approaches can be the starting point for future work investigating how to apply and the benefits of quantum machine learning approaches for data management tasks. In order to provide a solid basis for this future work, we will introduce the basics of quantum computing and quantum machine learning, and deliver a hands-on tutorial on applying quantum machine learning to one of the most important data management tasks—join order optimization.

Contributions: To the best of our knowledge, this is the first tutorial to discuss quantum machine learning approaches for DB research. This tutorial helps AI experts to get into quantum machine learning easily (by providing a comparison to classical ML), newbies to discover the possibilities of quantum computing by interacting with applications running on contemporary machines, and researchers to develop new methods utilizing quantum computing.

2 BASICS OF QUANTUM COMPUTING

For a database researcher applying machine learning, quantum computing appears as a promising model of computation. We start the tutorial by introducing the basic quantum computational concepts and viewing quantum computing as an extension of probabilistic computation to generalized notions of probability. The canonical introduction to quantum computing is Nielsen and Chuang's book [41], and the connection to probabilistic computing is based on Aaronson's approach [1] (also used in Ref. [33] and others).

Classical computation usually relies on bits `0` and `1`. Considering probabilistic computation, the outcome of an algorithm is not a single answer (`0` or `1`), but a probability distribution over the possible answers. Let $p_{\mathtt{0}}$ be the probability of observing bit `0`, and p_1 the probability of observing bit `1`. Since we must obtain some outcome, we have $p_{\mathtt{0}} + p_1 = 1$.

To formalize this scheme, we define that bits `0` and `1` are represented as vectors in the plane $\mathbb{R}^2$

$$|0\rangle := [1\ 0]^\top \qquad |1\rangle := [0\ 1]^\top . \tag{1}$$

The braket notation $|\cdot\rangle$ is commonly used in quantum computing to represent quantum states, and we (ab)use this notation to represent computational states. For single bits, a probability distribution over the possible bit states corresponds to a linear combination in $\mathbb{R}^2$

$$\vec{p} := p_{\mathtt{0}}\,|0\rangle + p_1\,|1\rangle$$

A probabilistic computation step needs to map vectors $\vec{p}$ to vectors that encode a valid probability distribution. Stochastic matrices S (i.e., the sum over columns is one for each column) are known to be the most general matrices to provide such a mapping. For instance, a NOT operation takes the form $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, a NOP (do nothing) operation is given by $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, and a probabilistic step that applies NOP or NOT with 50% chance is given by $\begin{pmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{pmatrix}$. A computational step is given by the transformation $\vec{p}' = S\vec{p}$.

Quantum mechanics (QM) exhibits a computational structure that is structurally similar to probabilistic computing. To "update" probabilistic to quantum computing (which should be more powerful than the former model), we need to translate the axioms of quantum mechanics into a list of requirements:

(1) Computations (*excluding* measurements that observe the current state of a computation) must be reversible.
(2) Vectors $\vec{p}$ comprise generalised probabilities over $\mathbb{C}$.
(3) Mapping between vectors $\vec{p}$ follows the dynamics of QM.

By changing to complex numbers (criterion 2), the probabilistic state vector $\vec{p}$, now written as general quantum state $|\varphi\rangle$, becomes

$$|\varphi\rangle := \alpha\,|0\rangle + \beta\,|1\rangle \tag{2}$$

$$\text{with } |\alpha|^2 + |\beta|^2 = 1, \tag{3}$$

where $\alpha, \beta \in \mathbb{C}$. $|\varphi\rangle$ is called a *qubit*, and is the fundamental unit of information in quantum computing.

Obviously, the complex coefficients α and β cannot be interpreted as probabilities any more, but quantum mechanics postulates that the probabilities of *observing* bit values 0 and 1, respectively, *after a measurement* of state $|\varphi\rangle$, are given by $p_0 = |\alpha|^2$ and $p_1 = |\beta|^2$. Essentially, this corresponds to switching from the 1-norm in the probabilistic picture to the 2-norm in quantum computing!

To ensure reversibility (criterion 1) and proper quantum mechanical dynamics (criterion 3), the most general matrix that transforms states as in Eq. (2) into other states, while obeying the normalization criterion in Eq. (3), is given by *unitary matrices.* A matrix U is unitary if its conjugate transpose matrix is its inverse [41].

The unitary matrices operating on qubits are called quantum logic gates or simply *gates.* In classical computing, the corresponding operations are logical operations such as AND and OR.

One of the most interesting and useful gates is the Hadamard gate $\hat{H}$, which is defined by the unitary matrix

$$\hat{H} := \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}.$$

The Hadamard gate does not have a corresponding classical operation. Applying the Hadamard gate to the basis vectors

$$\hat{H}\,|0\rangle = \frac{1}{\sqrt{2}}(|\,|0\rangle + |1\rangle \qquad \hat{H}\,|1\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) \tag{4}$$

shows that the gate maps the basis states $|0\rangle$ and $|1\rangle$ to superposition states where the probability to measure 0 or 1 is $1/2$.

Key differences between classical and quantum computing are superposition, entanglement, interference, and the role of measurements, which we all discuss in the tutorial.

Superposition states take the general form as given in Eq. (2); for a single qbit. This implies that a measurement delivers results 0 and 1 with non-zero probability. While such an outcome could also be obtained with classical probabilities, we will demonstrate in the tutorial how superpositions allow us to leverage complex generalized probabilities. The $\hat{H}$ gate creates a superposition state, as the calculations (4) show, leading to a stochastic measurement result. However, a second application $\hat{H}\hat{H}\,|0\rangle = 1/\sqrt{2}\hat{H}(|0\rangle + |1\rangle) = 1/2(|0\rangle + |1\rangle + |0\rangle - |1\rangle) = |0\rangle$ *removes* any stochasticity, which can be attributed to interference caused by the generalized probability of $-1/\sqrt{2}$ introduced when $\hat{H}$ is applied on $|1\rangle$.

Measurements are nonlinear operations. In the single qubit case, measuring means that we read either 0 or 1 depending on the probability distribution defined by the amplitudes of the state vector. For example, measuring the state in Eq. (2), we have 50% probability of measuring 0 or 1. This is not unlike sampling from a probability distribution. However, after measurement, a quantum state collapses,[1] leaving us with purely classical information.

Adiabatic quantum computing [4] is another quantum computing paradigm theoretically equivalent to universal circuit-based quantum computing [3, 66]. Adiabatic quantum computing is closely related to quantum annealing, which is implemented with quantum annealers. Although the current quantum annealers are not able to efficiently perform universal quantum computing [38], quantum annealing is a promising quantum computing method with a lot of industry-level applications [5, 6, 18, 20, 28, 40, 42, 43, 45–47, 51, 68]. We especially want to point out quantum annealing applications on scheduling transactions [7, 8], and multiple query optimization [60].

A single qubit quantum system has intuitive visualization with the so-called Bloch sphere. Because the probabilities are defined as

[1]Actually, the interpretation of measurements in quantum mechanics is subject of an intense ongoing physical and philosophical debate [52], but readers will forgive us if we stick to the orthodox Copenhagen interpretation for the sake of this tutorial.

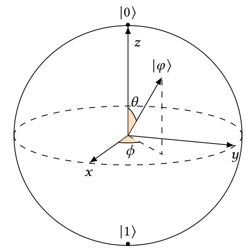

Figure 2: Bloch sphere

the lengths of the complex-valued amplitudes, we note that only the relative phase between the basis vectors $|0\rangle$ and $|1\rangle$ matters. Formally, because $|\alpha|^2 + |\beta|^2 = 1$, we can write the qubit

$$|\varphi\rangle = \alpha|0\rangle + \beta|1\rangle = e^{i\gamma}\left(\cos(\theta/2)|0\rangle + e^{\phi}\sin(\theta/2)|1\rangle\right),$$

where γ, θ, and ϕ are real-valued angles. A global phase $e^{i\gamma}$ cannot be observed in a measurement for physical reasons. Thus we can write a qubit using just two angles ϕ and θ. Now we can visualize the qubit as a point on the surface of the Bloch sphere. See Figure 2.

One of the most important gate sets is the rotation operators, which can be understood in the light of the Bloch sphere. They are especially useful in quantum machine learning because of their parametrization. For example, the rotation gate $R_x(\theta)$ is defined by the matrix

$$R_x(\theta) = \begin{bmatrix} \cos(\theta/2) & -i\sin(\theta/2) \\ -i\sin(\theta/2) & \cos(\theta/2) \end{bmatrix}.$$

The other rotation gates are $R_y(\theta)$ and $R_z(\theta)$ [41]. In the single qubit case, $R_x(\theta)$ (resp. $R_y(\theta)$ and $R_z(\theta)$) rotates the quantum state by an angle θ about the x (resp. y and z) axis of the Bloch sphere. Figure 4 shows an example of their usage.

2.1 Multi-qubit quantum states

A single qubit is the fundamental unit of quantum computation, but it is not useful alone. Quantum mechanical postulates define that we can construct multi-qubit quantum systems with the tensor product. Since qubits are elements of two-dimensional Hilbert space, the tensor product is a well-defined operation. For example, if we take the basis vectors of the single qubit quantum system, we obtain the computational basis as tensor products

$$|0\rangle \otimes |0\rangle = |00\rangle = [1\;0\;0\;0]^\top; \quad |0\rangle \otimes |1\rangle = |01\rangle = [0\;1\;0\;0]^\top$$
$$|1\rangle \otimes |0\rangle = |10\rangle = [0\;0\;1\;0]^\top; \quad |1\rangle \otimes |1\rangle = |11\rangle = [0\;0\;0\;1]^\top.$$

The dimension of the space grows exponentially as a function of n where n is the number of qubits. This exponential growth is among the reasons why QC is believed to perform better in certain tasks, resp. why a classical simulation of the underlying physics is hard.

As we mentioned before, entanglement is one of the key differences between classical and quantum computing. Entanglement means we cannot explain the entire system considering its pieces separately. In other words, the whole system is more than just the sum of its pieces. Entanglement appears in quantum computing when we have two or more qubits. The most common quantum logical gate to introduce entanglement between two qubits is the so-called controlled-NOT (CNOT) gate. The action of CNOT on the computational basis is defined by

$$|00\rangle \mapsto |00\rangle,\ |01\rangle \mapsto |01\rangle,\ |10\rangle \mapsto |11\rangle,\ |11\rangle \mapsto |10\rangle.$$

We see that the CNOT operation flips the second bit if the first bit is 1. Thus, the first bit controls the value of the second bit. Generally, the CNOT gate is defined by the unitary matrix

$$\text{CNOT} := \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

Universal quantum computations are usually visualized with circuit diagrams. See an example circuit in Figure 4.

There is a relatively small but standard set of quantum algorithms that can exhibit speedups over the corresponding best classical algorithms. The most famous algorithms are Shor's factoring algorithm [56], Grover's search algorithm [25], Deutsch-Jozsa algorithm [14], the quantum approximate optimization algorithm (QAOA) [17] and the quantum algorithm for solving semidefinite programs [9].

3 QUANTUM MACHINE LEARNING

For ease of understanding, we first introduce classical machine learning and extend it to quantum machine learning afterward. In machine learning, we want to learn to solve a problem from past experience instead of designing a new algorithm to solve the problem. For this, we want to create a model, which can predict the correct output for new input data based on past data. From the mathematical view, a model is a function $f(x, \theta)$, which calculates for a given input x the desired output $y = f(x, \theta)$ based on some parameters θ. Additionally, we have to define a quality measure $Q(x, y)$ for the generated output y of an input x. The goal of machine learning is to find the parameter vector θ that creates the best predicted outputs $\arg\max_{\theta} Q(x, f(x, \theta))$.

Multiple ML approaches exist, like supervised learning and reinforcement learning. In supervised learning, we train the model from given data points x_i with desired outputs y_i. Quality can be measured by comparing the predicted outputs $y = f(x_i, \theta)$ with the correct output y_i from the data. In reinforcement learning, the model controls the actions of an agent, which can interact with an environment by performing actions and receiving rewards. In contrast to supervised learning, we have no given data to learn from. Instead, the agent learns by interacting with the environment.

3.1 Hybrid quantum-classical algorithms

Quantum machine learning is commonly implemented as a hybrid quantum-classical algorithm [15], in which only a part of the calculation is performed on a quantum computer. A classical algorithm uses a quantum circuit as a function, passing classical inputs and receiving a classical result back. For quantum machine learning, the quantum circuit replaces the classical model. While the model is a quantum circuit, the optimization of the parameter vector θ is done by a classical algorithm. For this, a variational quantum circuit is used.

3.2 Variational quantum circuit

A variational quantum circuit (VQC), also called a parameterized quantum circuit, are quantum circuits that depend on some parameters. VQCs are proven to be universal approximators [48] like

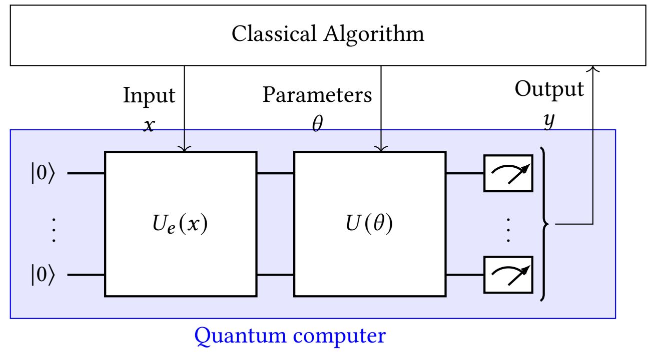

Figure 3: A hybrid algorithm using a VQC

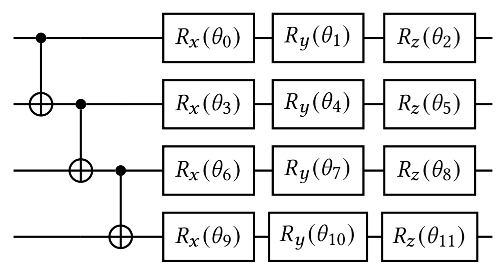

Figure 4: Example for a possible layer structure of a VQC

neural networks [29]. They are used in hybrid algorithms, where a classical algorithm optimizes the parameters over multiple runs to achieve the desired output.

A VQC always consists of three parts. The first part is the encoding layer, which encodes the classical data into a quantum state by applying a unitary operator $U_e(x)$ depending on the data. The second part is the calculation layer, which transforms the quantum state by applying a unitary operator $U(\theta)$ depending on parameters θ. The final part is the measurement layer, which obtains classical data from the quantum state. The classical algorithm interacts with the circuit by passing the input x as parameters to the encoding layer, passing the parameters θ to the calculation layers and receiving an output from the measurement layer (Fig 3).

The calculation layer is a unitary operator $U(\theta)$ that turns the quantum state representing the input into a quantum state representing the output. While any unitary operator is possible, it is common to use a combination of CNOT and rotation gates. The CNOT gates entangle the qubits to assure that every input qubit can affect every output qubit, and allows an advantage of a classical model by using the entanglement effect. The rotation gates are used as the parameterized operations.

A VQC commonly consists of repetitions of the same layer, which contains an entanglement and a rotation layer (Fig 4). While the layers use the same structure, they depend on different parameters and thus have different effects. The contribution [57] contains a comparison of the expressibility and entangling capabilities of different structures.

To optimize the parameters of the VQC, any optimizer can be used, but most often, a gradient descent optimizer is used. The gradients required for the optimization can be obtained by using the parameter shift rule [39]. Another possibility is to optimize the parameters using a genetic algorithm [11].

3.3 Encoding

To use quantum computing, we have to turn our classical data into a quantum state. The encoding is a unitary operator $U_e(x)$ that depends on input data x and turns the initial state into a quantum state representing our data. The choice of encoding influences the number of required qubits and the circuit depth. We introduce the three most common encoding methods [67] as follows.

(1) The simplest form is basis encoding, which encodes one classical bit into one qubit. This can be achieved by applying the X gate to a qubit if the corresponding classical bit is 1. This requires only one gate circuit depth, but only allows the encoding of binary data and requires many qubits. (2) Another encoding method is angle encoding, which uses the rotation gate to encode one real value into one qubit. This is done by using the input value as the parameter of a rotation gate. As the rotation gate is a periodical gate, the input has to be scaled to an interval smaller than 4π. This method allows a denser encoding than basis encoding and still requires only one gate circuit depth. (3) A denser encoding is possible using amplitude encoding, which encodes the values in the amplitudes of the quantum state. This allows the encoding of 2^n values into n qubits, with the limitation that the values have to be a vector of length 1. While amplitude encoding allows the densest encoding, it is also the most complex in circuit depth. The choice of the encoding method is a trade-off between the number of required qubits and the depth of the encoding circuit.

3.4 Decoding

To use a VQC in quantum machine learning, we have to obtain classical data from the circuit. The simplest solution is to use the output of the measurement. With n qubits, this provides a string of n bits, which represents one of the 2^n possible quantum states. This result is not deterministic, and the probability to obtain a result depends on the amplitude of the associated quantum state. The circuit can be run multiple times to approximate the probabilities of the different states. From these probabilities, we can choose the most likely state as our output. Alternatively, we can also use these probabilities as the outputs for our circuit, which provides us with a normalized vector of real values in the interval $[0, 1]$.

4 A DEMO ON QUANTUM MACHINE LEARNING WITH DATABASE APPLICATION

In our hands-on part of the tutorial, we will also present an implementation of a quantum machine learning application. Recent work [19, 31, 58] provides promising initial results for quantum machine learning in general synthetic environments. Likewise, machine learning has become a popular approach to solving various problems that occur in data management systems, for instance, join order optimization [37, 61], indexing [36] or database tuning [35]. Especially given that quantum speedups on noisy machines have been shown on solid theoretical grounds [27], which is a feat not achieved by most known NISQ algorithms, combining the two approaches seems to be a natural match.

Since the combination of quantum machine learning and database problems is still in its very infancy, with an extremely small amount of published papers, we focus on an initial setup of a quantum machine learning algorithm for database tasks. To this end, we will present a novel quantum-assisted machine learning technique for join order optimization. The basic idea of our algorithm is to learn a reward function that represents the join order quality with a VQC as a machine learning model. This demonstration intends to enable database researchers to understand the general principles well enough to start exploring how to couple the ideas with their domains of research.

The implementation will be in Python using the machine learning library *PyTorch* [44]. We will discuss implementation details of the VQC definition (see Section 3.2) and data encoding (see Section 3.3) for the join order problem. These quantum-based parts use the quantum framework *Qiskit* [59]. In the end, we will train a (simulated) quantum machine learning model and compare the join order quality with existing optimizers.

5 OPPORTUNITIES AND FUTURE WORK FOR DATABASE RESEARCH

5.1 State-of-the-art

There already exists some work about applying quantum computing to solve database tasks, including multiple query optimization [16, 60], transaction scheduling [7, 8, 23], and join order optimization [53, 55, 63]. Multiple Query Optimization (MQO) is an important NP-hard problem in databases. [60] first tackles the MQO problem with the D-Wave adiabatic quantum annealer. Then [16] use a different quantum computer to solve the MQO problem with a hybrid classical-quantum algorithm on a gate-based quantum computer. Although the problem size is limited due to the limitations of current quantum devices, the experimental results of [60] and [16] have already shown some advantages compared with the classical solutions. Transaction scheduling problem aims to determine the optimal order of parallel execution of transactions for best performance. In [7, 8, 23], the authors transform an instance of the transaction schedule problem into a formula that is accepted by quantum annealers. Experimental evaluation shows the runtime on a quantum annealer outperforms the runtime of traditional algorithms. Join order selection is an NP-hard problem in relational databases. [53, 55] formulate the join order selection as a quadratic unconstrained binary optimization (QUBO) problem and solved it on two state-of-the-art approaches (gate-based quantum computing and quantum annealing). The DB-QPU co-design approach is proposed to overcome the limitations of current quantum devices. It should be noticed that most of the previous works of quantum computing for databases are not about quantum machine learning. How quantum machine learning could benefit classical machine learning for databases is still an open problem.

5.2 Open Challenges

While some progress has already been made, the research on quantum computing for databases has just begun, and there are many opportunities and open problems for further exploration.

5.2.1 From the database perspective. Machine learning techniques have been proposed to optimize data management in recent years. Theoretically, all the existing works that utilize classical machine learning for databases, could be enhanced or replaced with their quantum counterparts. For example, we could use quantum neural networks for cardinality estimation and quantum reinforcement learning for join order selection. Both the online and offline optimization tasks could be solved with QML methods. Online tasks of query optimization include cardinality estimation, cost estimation, execution time prediction, resource utilization estimation, query rewrite, join order selection, query scheduling, and transaction scheduling. Offline optimization tasks include knob tuning, index selection, materialized view recommendation, and data partition.

5.2.2 From the quantum perspective. From the perspective view of quantum computing and quantum machine learning, some basic problems also need to be further researched, such as: (1) What are the advantages of using quantum machine learning for databases compared to classical machine learning? Are the quantum machine learning algorithms really suitable for replacing their classical counterparts for databases? (2) How to combine classical and quantum algorithms to achieve good speedups with few qubits? (3) How to improve the scalability of QC algorithms in the near-term quantum computers since the databases usually need to cope with large sizes of queries? (4) How (and if) can payload data be included?

6 TUTORIAL ORGANIZATION

The tutorial is planned for 3 hours (180 minutes) and will have the following structure:

I. Introduction and motivation (10'). We introduce the background and remarkable progress of quantum computing. We show a correlation exists between the availability of larger quantum computers and the publication performances of researchers in the areas of quantum computing and quantum machine learning.

II. Basics of Quantum Computing (50'). We present the basics of quantum computing techniques, including quantum bits, Bloch sphere, multi-qubit states.

III. Quantum machine learning (50'). We introduce quantum machine learning techniques, including hybrid quantum-classical algorithms, variational quantum circuits, encoding, and decoding.

IV. Break and QA (20') We allocate time to answer the questions and encourage interaction with the audience.

V. Demo about join order optimization with the quantum machine learning (30') We design a demo to demonstrate how to use quantum machine learning to perform join order optimization with the framework Qiskit [59].

VI. Open problems and challenges for database research (20'). We discuss the state-of-the-art and open challenges to applying quantum machine learning for database research.

7 GOAL OF THIS TUTORIAL

Intended Audience. This tutorial is intended for a wide scope of audience ranging from academic researchers and students to industrial developers and practitioners that want to understand the impact of quantum machine learning on databases. Basic knowledge of linear algebra and the quantum mechanism is sufficient to follow the tutorial. Some background in machine learning and simulated annealing algorithms would be useful.

Learning Outcomes. The main learning outcomes of this tutorial are: (1) understanding the impact and remarkable progress of quantum computing; (2) learning basics on quantum computing, such as superposition, entanglement, Bloch sphere, and multi-qubit quantum states; (3) learning the basics of quantum machine learning, including variational quantum circuits, encoding, and decoding data. (4) identifying open problems and research challenges of quantum machine learning for databases. Practitioners and students will be able to quickly build an extensive understanding as well as grasp the latest trends and state-of-the-art techniques in quantum machine learning. In addition, this tutorial will provide a hands-on demonstration to guide both researchers and developers.

8 SHORT BIOGRAPHIES

Tobias Winker is a Ph.D. student at the University of Lübeck with a master's degree in computer science. He is a member of the QC4DB project. His research interests are classical and quantum machine learning for database problems [24].

Sven Groppe is a Professor at the University of Lübeck and the project coordinator of the QC4DB project. His research includes the integration [21, 22] of quantum accelerators into DBMS [7, 8, 23, 24], and high-level quantum programming languages [26]. He is a chair of SBD@SIGMOD (2016-2020), BiDEDE@SIGMOD (2021-2023), VLIoT@VLDB (2017-2022) and QDSM@VLDB 2023.

Valter Uotila is a Ph.D. student at the University of Helsinki. His research interests are in the intersection of databases, quantum computing, and category theory [63–65].

Zhengtong Yan is a Ph.D. student at the University of Helsinki. His research topics lie in quantum computing and reinforcement learning for databases.

Jiaheng Lu is a Professor at the University of Helsinki. His current research interests focus on multi-model databases and quantum computing for database applications. He has written four books on Hadoop and NoSQL databases, and more than 130 journal and conference papers published in SIGMOD, VLDB, TODS, etc.

Maja Franz is a research master student at the Technical University of Applied Sciences Regensburg. Her research interests focus on quantum algorithms for near-term quantum devices [19, 54], especially in industrially relevant domains. After her master's she intends to obtain a Ph.D. degree in the field of quantum computing.

Wolfgang Mauerer is a Professor at the Technical University of Applied Sciences Regensburg, and a Senior Research Scientist at Siemens Technology. His interests focus on software/systems engineering, and quantum computing. He has published strongly multidisciplinary work in venues and journals from Nature Photonics and PRL via ICSE and TSE to SIGMOD.

ACKNOWLEDGMENTS

This work is funded by the German Federal Ministry of Education and Research within the funding program quantum technologies—from basic research to market—contract numbers 13N16090 (TW, SG), 13N15645 (WM) and 13N16092 (MF). WM and MF also acknowledge support from the High-Tech Agenda Bavaria.

REFERENCES

[1] Scott Aaronson. 2013. *Quantum Computing since Democritus.* Cambridge University Press, USA.

[2] Zainab Abohashima, Mohamed Elhosen, Essam H. Houssein, and Waleed M. Mohamed. 2020. Classification with Quantum Machine Learning: A Survey. *arXiv* arXiv:2006.12270 (2020). https://doi.org/10.48550/ARXIV.2006.12270

[3] Dorit Aharonov, Wim van Dam, Julia Kempe, Zeph Landau, Seth Lloyd, and Oded Regev. 2007. Adiabatic Quantum Computation is Equivalent to Standard Quantum Computation. *SIAM J. Comput.* 37, 1 (2007), 166–194. https://doi.org/10.1137/S0097539705447323 arXiv:https://doi.org/10.1137/S0097539705447323

[4] T. Albash and D.A. Lidar. 2018. Adiabatic quantum computation. *Reviews of Modern Physics* 90, 1 (2018). https://doi.org/10.1103/RevModPhys.90.015002

[5] Ryan Babbush, Alejandro Perdomo-Ortiz, Bryan O'Gorman, William Macready, and Alan Aspuru-Guzik. 2014. *Construction of Energy Functions for Lattice Heteropolymer Models: Efficient Encodings for Constraint Satisfaction Programming and Quantum Annealing.* John Wiley & Sons, Inc., Hoboken, New Jersey, 201–244. https://doi.org/10.1002/9781118755815.ch05

[6] Zhengbing Bian, Fabian Chudak, Robert Brian Israel, Brad Lackey, William G. Macready, and Aidan Roy. 2016. Mapping Constrained Optimization Problems to Quantum Annealing with Application to Fault Diagnosis. *Frontiers in ICT* 3 (Jul 2016). https://doi.org/10.3389/fict.2016.00014

[7] Tim Bittner and Sven Groppe. 2020. Avoiding blocking by scheduling transactions using quantum annealing. In *Proceedings of the 24th Symposium on International Database Engineering & Applications.* ACM. https://doi.org/10.1145/3410566.3410593

[8] Tim Bittner and Sven Groppe. 2020. Hardware Accelerating the Optimization of Transaction Schedules via Quantum Annealing by Avoiding Blocking. *Open Journal of Cloud Computing (OJCC)* 7, 1 (2020), 1–21. http://nbn-resolving.de/urn:nbn:de:101:1-2020112218332015343957

[9] Fernando G. S. L. Brandao and Krysta Svore. 2016. Quantum Speed-ups for Semidefinite Programming. (2016). https://doi.org/10.48550/ARXIV.1609.05537

[10] Matthias C Caro, Hsin-Yuan Huang, M Cerezo, Kunal Sharma, Andrew Sornborger, Lukasz Cincio, and Patrick J Coles. 2022. Generalization in quantum machine learning from few training data. *Nat. Commun.* 13, 1 (2022).

[11] Samuel Yen-Chi Chen, Chih-Min Huang, Chia-Wei Hsing, Hsi-Sheng Goan, and Ying-Jer Kao. 2022. Variational quantum reinforcement learning via evolutionary optimization. *Machine Learning: Science and Technology* 3, 1 (2 2022), 015025. https://doi.org/10.1088/2632-2153/ac4559

[12] Carlo Ciliberto, Mark Herbster, Alessandro Davide Ialongo, Massimiliano Pontil, Andrea Rocchetto, Simone Severini, and Leonard Wossnig. 2018. Quantum machine learning: a classical perspective. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 474, 2209 (Jan. 2018), 20170551. https://doi.org/10.1098/rspa.2017.0551

[13] Andrew W. Cross, Lev S. Bishop, Sarah Sheldon, Paul D. Nation, and Jay M. Gambetta. 2019. Validating quantum computers using randomized model circuits. *Physical Review A* 100, 3 (Sept. 2019). https://doi.org/10.1103/physreva.100.032328

[14] David Deutsch and Richard Jozsa. 1992. Rapid solution of problems by quantum computation. *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences* 439, 1907 (Dec 1992), 553–558. https://doi.org/10.1098/rspa.1992.0167

[15] Suguru Endo, Zhenyu Cai, Simon C. Benjamin, and Xiao Yuan. 2021. Hybrid Quantum-Classical Algorithms and Quantum Error Mitigation. *Journal of the Physical Society of Japan* 90, 3 (mar 2021), 032001. https://doi.org/10.7566/jpsj.90.032001

[16] Tobias Fankhauser, Marc E Solèr, Rudolf M Füchslin, and Kurt Stockinger. 2021. Multiple query optimization using a hybrid approach of classical and quantum computing. *arXiv preprint arXiv:2107.10508* (2021).

[17] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. 2014. A Quantum Approximate Optimization Algorithm. (2014). https://doi.org/10.48550/ARXIV.1411.4028

[18] Marina Fernández-Campoamor, Corey O'Meara, Giorgio Cortiana, Vedran Peric, and Juan Bernabé-Moreno. 2021. Community Detection in Electrical Grids Using Quantum Annealing. arXiv:2112.08300 (Dec 2021). https://doi.org/10.48550/arXiv.2112.08300 arXiv:2112.08300 [quant-ph].

[19] Maja Franz, Lucas Wolf, Maniraman Periyasamy, Christian Ufrecht, Daniel D. Scherer, Axel Plinge, Christopher Mutschler, and Wolfgang Mauerer. 2022. Uncovering instabilities in variational-quantum deep Q-networks. *Journal of the Franklin Institute* (2022). https://doi.org/10.1016/j.jfranklin.2022.08.021

[20] Erica Grant, Travis S. Humble, and Benjamin Stump. 2021. Benchmarking Quantum Annealing Controls with Portfolio Optimization. *Physical Review Applied* 15, 1 (Jan 2021), 014012. https://doi.org/10.1103/PhysRevApplied.15.014012

[21] Sven Groppe. 2021. Semantic Hybrid Multi-Model Multi-Platform (SHM3P) Databases. In *International Semantic Intelligence Conference (ISIC 2021), New Delhi (hybrid), India.* CEUR, 16–26. http://ceur-ws.org/Vol-2786/Paper2.pdf

[22] Sven Groppe and Jinghua Groppe. 2020. Hybrid Multi-Model Multi-Platform (HM3P) Databases. In *Proceedings of the 9th International Conference on Data Science, Technology and Applications (DATA).* https://doi.org/10.5220/0009802401770184

[23] Sven Groppe and Jinghua Groppe. 2021. Optimizing Transaction Schedules on Universal Quantum Computers via Code Generation for Grover's Search Algorithm. In *25th International Database Engineering & Applications Symposium.* ACM. https://doi.org/10.1145/3472163.3472164

[24] Sven Groppe, Jinghua Groppe, Umut Çalıkyılmaz, Tobias Winker, and Le Gruenwald. 2022. Quantum Data Management and Quantum Machine Learning for Data Management: State-of-the-Art and Open Challenges. In *Proceedings of the EAI ICISML conference.*

[25] Lov K. Grover. 1996. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing - STOC '96.* ACM Press, Philadelphia, Pennsylvania, United States, 212–219. https://doi.org/10.1145/237814.237866

[26] Julian Hans and Sven Groppe. 2022. Silq2Qiskit - Developing a quantum language source-to-source translator. In *Proceedings of the 5th International Conference on Computer Science and Software Engineering (CSSE 2022)* (Guilin, China). https://doi.org/10.1145/3569966.3570114

[27] Vojtěch Havlíček, Antonio D. Córcoles, Kristan Temme, Aram W. Harrow, Abhinav Kandala, Jerry M. Chow, and Jay M. Gambetta. 2019. Supervised learning with quantum-enhanced feature spaces. *Nature* 567, 7747 (2019), 209–212. https://doi.org/10.1038/s41586-019-0980-2

[28] Itay Hen and A. P. Young. 2012. Solving the graph-isomorphism problem with a quantum annealer. *Physical Review A* 86, 4 (Oct 2012), 042310. https://doi.org/10.1103/PhysRevA.86.042310

[29] K. Hornik, M. Stinchcombe, and H. White. 1989. Multilayer feedforward networks are universal approximators. *Neural Networks* 2, 5 (1989), 359–366.

[30] 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. *Nat. Commun.* 12, 1 (2021).

[31] Sofiene Jerbi, Casper Gyurik, Simon C. Marshall, Hans J. Briegel, and Vedran Dunjko. 2021. Parametrized quantum policies for reinforcement learning. https://doi.org/10.48550/ARXIV.2103.05577

[32] Arun Kumar, Matthias Boehm, and Jun Yang. 2017. Data Management in Machine Learning: Challenges, Techniques, and Systems. In *Proceedings of the 2017 ACM International Conference on Management of Data* (Chicago, Illinois, USA) *(SIGMOD '17).* Association for Computing Machinery, New York, NY, USA, 1717–1722. https://doi.org/10.1145/3035918.3054775

[33] Joseph M. Landsberg. 2019. *A Very Brief Introduction to Quantum Computing and Quantum Information Theory for Mathematicians.* Springer International Publishing, Cham, 5–41. https://doi.org/10.1007/978-3-030-06122-7_2

[34] Guoliang Li, Xuanhe Zhou, and Lei Cao. 2021. AI meets database: AI4DB and DB4AI. In *Proceedings of the 2021 International Conference on Management of Data.* 2859–2866.

[35] Guoliang Li, Xuanhe Zhou, Shifu Li, and Bo Gao. 2019. QTune: A Query-Aware Database Tuning System with Deep Reinforcement Learning. *Proc. VLDB Endow.* 12, 12 (aug 2019), 2118–2130. https://doi.org/10.14778/3352063.3352129

[36] Gabriel Paludo Licks and Felipe Meneguzzi. 2020. Automated Database Indexing using Model-free Reinforcement Learning. https://doi.org/10.48550/ARXIV.2007.14244

[37] Ryan Marcus and Olga Papaemmanouil. 2018. Deep Reinforcement Learning for Join Order Enumeration. In *Proceedings of the First International Workshop on Exploiting Artificial Intelligence Techniques for Data Management* (Houston, TX, USA) *(aiDM'18).* Association for Computing Machinery, New York, NY, USA, Article 3, 4 pages. https://doi.org/10.1145/3211954.3211957

[38] Catherine C. McGeoch. 2020. Theory versus practice in annealing-based quantum computing. *Theoretical Computer Science* 816 (May 2020), 169–183. https://doi.org/10.1016/j.tcs.2020.01.024

[39] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii. 2018. Quantum circuit learning. *Physical Review A* 98, 3 (sep 2018). https://doi.org/10.1103/physreva.98.032309

[40] Florian Neukart, Gabriele Compostella, Christian Seidel, David von Dollen, Sheir Yarkoni, and Bob Parney. 2017. Traffic Flow Optimization Using a Quantum Annealer. *Frontiers in ICT* 4 (Dec 2017), 29. https://doi.org/10.3389/fict.2017.00029

[41] Michael A. Nielsen and Isaac L. Chuang. 2010. *Quantum computation and quantum information* (10th anniversary ed ed.). Cambridge University Press, Cambridge, New York.

[42] Masayuki Ohzeki, Akira Miki, Masamichi J. Miyama, and Masayoshi Terabe. 2019. Control of Automated Guided Vehicles Without Collision by Quantum Annealer and Digital Devices. *Frontiers in Computer Science* 1 (Nov 2019), 9. https://doi.org/10.3389/fcomp.2019.00009

[43] Daniel O'Malley, Velimir V. Vesselinov, Boian S. Alexandrov, and Ludmil B. Alexandrov. 2018. Nonnegative/Binary matrix factorization with a D-Wave quantum annealer. *PLOS ONE* 13, 12 (Dec 2018), e0206653. https://doi.org/10.1371/journal.pone.0206653

[44] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32.* Curran

Associates, Inc., 8024–8035. http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf

[45] WangChun Peng, BaoNan Wang, Feng Hu, YunJiang Wang, XianJin Fang, XingYuan Chen, and Chao Wang. 2019. Factoring larger integers with fewer qubits via quantum annealing with optimized parameters. *Science China Physics, Mechanics & Astronomy* 62, 6 (Jun 2019), 60311. https://doi.org/10.1007/s11433-018-9307-1

[46] Alejandro Perdomo-Ortiz, Neil Dickson, Marshall Drew-Brook, Geordie Rose, and Alán Aspuru-Guzik. 2012. Finding low-energy conformations of lattice protein models by quantum annealing. *Scientific Reports* 2, 1 (Dec 2012), 571. https://doi.org/10.1038/srep00571

[47] A. Perdomo-Ortiz, J. Fluegemann, S. Narasimhan, R. Biswas, and V.N. Smelyanskiy. 2015. A quantum annealing approach for fault detection and diagnosis of graph-based systems. *The European Physical Journal Special Topics* 224, 1 (Feb 2015), 131–148. https://doi.org/10.1140/epjst/e2015-02347-y

[48] Adrián Pérez-Salinas, Alba Cervera-Lierta, Elies Gil-Fuster, and José I. Latorre. 2020. Data re-uploading for a universal quantum classifier. *Quantum* 4 (Feb. 2020), 226. https://doi.org/10.22331/q-2020-02-06-226

[49] Neoklis Polyzotis, Sudip Roy, Steven Euijong Whang, and Martin Zinkevich. 2017. Data Management Challenges in Production Machine Learning. In *Proceedings of the 2017 ACM International Conference on Management of Data* (Chicago, Illinois, USA) *(SIGMOD '17)*. Association for Computing Machinery, New York, NY, USA, 1723–1726. https://doi.org/10.1145/3035918.3054782

[50] Patrick Rebentrost, Masoud Mohseni, and Seth Lloyd. 2014. Quantum Support Vector Machine for Big Data Classification. *Phys. Rev. Lett.* 113 (2014), 130503. Issue 13. https://doi.org/10.1103/PhysRevLett.113.130503

[51] Eleanor G. Rieffel, Davide Venturelli, Bryan O'Gorman, Minh B. Do, Elicia M. Prystay, and Vadim N. Smelyanskiy. 2015. A case study in programming a quantum annealer for hard operational planning problems. *Quantum Information Processing* 14, 1 (Jan 2015), 1–36. https://doi.org/10.1007/s11128-014-0892-x

[52] Maximilian Schlosshauer. 2005. Decoherence, the measurement problem, and interpretations of quantum mechanics. *Rev. Mod. Phys.* 76 (Feb 2005), 1267–1305. Issue 4. https://doi.org/10.1103/RevModPhys.76.1267

[53] Manuel Schönberger. 2022. Applicability of Quantum Computing on Database Query Optimization. In *Proceedings of the 2022 International Conference on Management of Data (SIGMOD), Philadelphia, PA, USA (SIGMOD '22)*. Association for Computing Machinery, New York, NY, USA, 2512–2514. https://doi.org/10.1145/3514221.3520257

[54] Manuel Schönberger, Maja Franz, Stefanie Scherzinger, and Wolfgang Mauerer. 2022. Peel | Pile? Cross-Framework Portability of Quantum Software. *QSA@IEEE International Conference on Software Architecture (ICSA)* (2022). https://doi.org/10.1109/ICSA-C54293.2022.00039

[55] Manuel Schönberger, Stefanie Scherzinger, and Wolfgang Mauerer. 2023. Ready to Leap (by Co-Design)? Join Order Optimisation on Quantum Hardware. In *Proceedings of ACM SIGMOD/PODS International Conference on Management of Data*.

[56] Peter W. Shor. 1997. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM J. Comput.* 26, 5 (Oct 1997), 1484–1509. https://doi.org/10.1137/S0097539795293172

[57] Sukin Sim, Peter D. Johnson, and Alá n Aspuru-Guzik. 2019. Expressibility and Entangling Capability of Parameterized Quantum Circuits for Hybrid Quantum-Classical Algorithms. *Advanced Quantum Technologies* 2, 12 (2019), 1900070. https://doi.org/10.1002/qute.201900070

[58] Andrea Skolik, Sofiene Jerbi, and Vedran Dunjko. 2022. Quantum agents in the gym: a variational quantum algorithm for deep q-learning. *Quantum* 6 (2022), 720. https://doi.org/10.22331/q-2022-05-24-720

[59] Matthew Treinish, Jay Gambetta, Paul Nation, Paul Kassebaum, Diego M. Rodríguez, Salvador de la Puente González, Jake Lishman, Shaohan Hu, Kevin Krsulich, Jim Garrison, Luciano Bello, Jessie Yu, Manoel Marques, Julien Gacon, David McKay, Juan Gomez, Lauren Capelluto, Travis-S-IBM, Abby Mitchell, Ashish Panigrahi, lerongil, Rafey Iqbal Rahman, Steve Wood, Toshinari Itoko, Alex Pozas-Kerstjens, Christopher J. Wood, Divyanshu Singh, Drew Risinger, and Eli Arbel. 2022. *Qiskit*. https://doi.org/10.5281/zenodo.7416349

[60] Immanuel Trummer and Christoph Koch. 2016. Multiple query optimization on the D-Wave 2X adiabatic quantum computer. *Proceedings of the VLDB Endowment* 9, 9 (May 2016), 648–659. https://doi.org/10.14778/2947618.2947621

[61] Immanuel Trummer, Junxiong Wang, Ziyun Wei, Deepak Maram, Samuel Moseley, Saehan Jo, Joseph Antonakakis, and Ankush Rayabhari. 2021. SkinnerDB: Regret-Bounded Query Evaluation via Reinforcement Learning. *ACM Trans. Database Syst.* 46, 3, Article 9 (sep 2021), 45 pages. https://doi.org/10.1145/3464389

[62] Fatma undefinedzcan, Abdul Quamar, Jaydeep Sen, Chuan Lei, and Vasilis Efthymiou. 2020. State of the Art and Open Challenges in Natural Language Interfaces to Data. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data* (Portland, OR, USA) *(SIGMOD '20)*. Association for Computing Machinery, New York, NY, USA, 2629–2636. https://doi.org/10.1145/3318464.3383128

[63] Valter Uotila. 2022. Synergy between Quantum Computers and Databases. In *Proceedings of the VLDB 2022 PhD Workshop co-located with the 48th International Conference on Very Large Databases (VLDB 2022), Sydney, Australia (CEUR Workshop Proceedings, Vol. 3186)*, Zhifeng Bao and Timos K. Sellis (Eds.). CEUR-WS.org. http://ceur-ws.org/Vol-3186/paper_1.pdf

[64] Valter Uotila and Jiaheng Lu. 2021. A Formal Category Theoretical Framework for Multi-model Data Transformations. In *Heterogeneous Data Management, Polystores, and Analytics for Healthcare*, El Kindi Rezig, Vijay Gadepally, Timothy Mattson, Michael Stonebraker, Tim Kraska, Fusheng Wang, Gang Luo, Jun Kong, and Alevtina Dubovitskaya (Eds.). Springer International Publishing, Cham, 14–28.

[65] Valter Uotila, Jiaheng Lu, Dieter Gawlick, Zhen Hua Liu, Souripriya Das, and Gregory Pogossiants. 2021. MultiCategory: Multi-Model Query Processing Meets Category Theory and Functional Programming. *Proc. VLDB Endow.* 14, 12 (jul 2021), 2663–2666. https://doi.org/10.14778/3476311.3476314

[66] W. van Dam, M. Mosca, and U. Vazirani. 2001. How powerful is adiabatic quantum computation?. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*. IEEE, Newport Beach, CA, USA, 279–287. https://doi.org/10.1109/SFCS.2001.959902

[67] Manuela Weigold, Johanna Barzen, Frank Leymann, and Marie Salm. 2021. Encoding patterns for quantum algorithms. *IET Quantum Communication* 2, 4 (2021), 141–152. https://doi.org/10.1049/qtc2.12032

[68] Sheir Yarkoni, Alex Alekseyenko, Michael Streif, David Von Dollen, Florian Neukart, and Thomas Back. 2021. Multi-car paint shop optimization with quantum annealing. In *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE, Broomfield, CO, USA, 35–41. https://doi.org/10.1109/QCE52317.2021.00019

[69] Renxin Zhao and Shi Wang. 2021. A review of Quantum Neural Networks: Methods, Models, Dilemma. *arXiv* 2109.01840 (2021). https://doi.org/10.48550/ARXIV.2109.01840