

# Relevance of Near-Term Quantum Computing in the Cloud: A Humanities Perspective

Johanna Barzen, Frank Leymann, Michael Falkenthal, Daniel Vietz, Benjamin Weder, Karoline Wild

> Institute of Architecture of Application Systems, University of Stuttgart, Germany {Johanna.Barzen}@iaas.uni-stuttgart.de

# ВівТ<sub>Е</sub>Х

@article {Barzen2021 QuantumComputing, Author = {Johanna Barzen and Frank Leymann and Michael Falkenthal and Daniel Vietz and Benjamin Weder and Karoline Wild}, Title = {{Relevance of Near-Term Quantum Computing in the Cloud: A Humanities Perspective}}, Journal = {Cloud Computing and Services Science}, Publisher = {Springer International Publishing}, Volume = 1399, Pages = {25--58}, Month = mar, Year = 2021, Url = {https://link.springer.com/chapter/10.1007%2F978-3-030-72369-9\_2}, }

1

© Springer Nature Switzerland AG 2021

This is a post-peer-review, pre-copyedit version of an article published in Ferguson D., Pahl C., Helfert M. (eds) Cloud Computing and Services Science. CLOSER 2020. Communications in Computer and Information Science, vol 1399, Springer, Cham. The final authenticated version is available online at: https://doi.org/10.1007/978-3-030-72369-9\_2



Universität Stuttgart Germany

# **Relevance of Near-Term Quantum Computing** in the Cloud: A Humanities Perspective

Johanna Barzen<sup>1</sup>[0000-0001-8397-7973]</sup>, Frank Leymann<sup>1</sup>[0000-0002-9123-259X], Michael Falkenthal<sup>2</sup>[0000-0001-7802-1395]</sup>, Daniel Vietz<sup>1</sup>[0000-0003-1366-5805]</sup>, Benjamin Weder<sup>1</sup>[0000-0002-6761-6243]</sup>, and Karoline Wild<sup>1</sup>[0000-0001-7803-6386]

> <sup>1</sup> University of Stuttgart, Universitätsstraße 38, 70569 Stuttgart, Germany {firstname.lastname}@iaas.uni-stuttgart.de <sup>2</sup> StoneOne AG, Keithstraße 6, 10787 Berlin, Germany michael.falkenthal@stoneone.de

**Abstract.** As quantum computers are becoming real, they have the inherent potential to significantly impact many application domains. In this paper we outline the fundamentals about programming quantum computers and show that quantum programs are typically hybrid consisting of a mixture of classical parts and quantum parts. With the advent of quantum computers in the cloud, the cloud is a fine environment for performing quantum programs. The tool chain available for creating and running such programs is sketched. As an exemplary problem we discuss efforts to implement quantum programs that are hardware independent. A use case from quantum humanities is discussed, hinting which applications in this domain can already be used in the field of (quantum) machine learning. Finally, a collaborative platform for solving problems with quantum computers – that is currently under construction – is presented.

**Keywords:** Cloud Computing, Quantum Computing, Hybrid Applications, Quantum Humanities.

## 1 Introduction

Quantum computing advanced up to a state that urges attention to the software community: problems that are hard to solve based on classical (hardware and software) technology become tractable in the next couple of years [54]. Quantum computers are offered for commercial use (e.g. IBM Q System One), and access to quantum computers are offered by various vendors like Amazon, IBM, Microsoft, or Rigetti via the cloud.

However, todays quantum computers are error-prone. For example, the states they store are volatile and decay fast (decoherence), the operations they perform are not exact (gate fidelity) etc. Consequently, they are "noisy". And their size (measured in Qubits – see section 2.1) is of "intermediate scale". Together, todays quantum computers are Noisy Intermediate Scale Quantum (NISQ) computers [61]. In order to perform a quantum algorithm reliably on a NISQ machine, it must be limited in size [41].

Because of this, the overall algorithms are often hybrid. They perform parts on a quantum computer, other parts on a classical computer. Each part performed on a quantum computer is fast enough to produce reliable results. The parts executed on a classical computer analyze the results, compute new parameters for the quantum parts, and pass them on to a quantum part. Typically, this is an iteration consisting of classical pre-processing, quantum processing, and classical post-processing.

This iteration between classical parts and quantum parts reveals why the cloud is a solid basis for executing quantum applications: it offers classical environments as well as quantum computers (see before).

What are viable applications on NISQ computers? For example, simulation of molecules in drug discovery or material science is very promising [26], many areas of machine learning will realize significant improvements [18], as well as solving optimization problems [27].

The reminder of this paper is structured as following: Section 2 sketches the programming model of quantum computers. Quantum computing in the cloud is introduced in section 3. How to remove hardware dependencies is addressed in section 4. Section 5 outlines a use case from quantum humanities. A collaboration platform for developing and exploiting quantum applications is subject of section 6. Section 7 concludes the paper.

This paper is based on [44] extended by a comprehensive use case for quantum machine learning from the digital humanities.

# 2 Programming Model

Next, we introduce the basics of the quantum programming model – see [55].

#### 2.1 Quantum Registers

The most fundamental notion of quantum computing is the quantum bit or qubit for short. While a classical bit can have either the value 0 or 1 at a given time, the value of a qubit  $|x\rangle$  is any combination of these two values:  $|x\rangle = \alpha \cdot |0\rangle + \beta \cdot |1\rangle$  (to distinguish bits from qubits we write  $|x\rangle$  instead of x for the latter). This so-called superposition is one source of the power of quantum computing.

The actual value of a qubit is determined by a so-called measurement.  $|\alpha|^2$  and  $|\beta|^2$  are the probabilities that – once the qubit is measured – the classical value "0" or "1", respectively, results. Because either "0" or "1" will definitively result, the probabilities sum up to 1:  $|\alpha|^2 + |\beta|^2 = 1$ .

Just like bits are combined into registers in a classical computer, qubits are combined into quantum registers. A quantum register  $|r\rangle$  consisting of n qubits has a value that is a superposition of the 2<sup>n</sup> values  $|0...0\rangle$ ,  $|0...01\rangle$ , up to  $|1...1\rangle$ . A manipulation of the quantum register thus modifies these 2<sup>n</sup> values at the same time: this quantum parallelism is another source of the power of quantum computing.

#### 2.2 Quantum Operations

Figure 1 depicts two qubits  $\alpha|0\rangle+\beta|1\rangle$  and  $\gamma|0\rangle+\delta|1\rangle$ : because  $|\alpha|^2+|\beta|^2 = |\gamma|^2+|\delta|^2 = 1$ , each qubit can be represented as a point on the unit circle, i.e. as a vector of length 1. Manipulating a qubit results in another qubit, i.e. a manipulation U of qubits preserves the lengths of qubits as vectors. Such manipulations are called unitary transformations. A quantum algorithm combines such unitary transformations to manipulate qubits (or quantum registers in general). Since the combination of unitary transformations is again a unitary transformation, a quantum algorithm is represented by a unitary transformation.



Fig. 1. Depicting a qubit and its manipulation (see [44] Fig. 1)

This geometric interpretation of qubits is extended to quantum registers: a quantum register with n qubits can be perceived as a unit vector in a  $2^n$ -dimensional vector space. A quantum algorithm is then a unitary transformation of this vector space.

A quantum algorithm U takes a quantum register  $|r\rangle$  as input and produces a quantum register  $|s\rangle=U(|r\rangle)$  as output, with

$$|s\rangle = \sum_{i=1}^{2^n} \alpha_i |x_1^i \cdots x_n^i\rangle, x_j^i \in \{0,1\}$$

The actual result of the algorithm U is determined by measuring  $|s\rangle$ . Thus, the result is  $(x_1^i \cdots x_n^i) \in \{0,1\}^n$  with probability  $\alpha_i^2$ . Obviously, different executions of U followed by a measurement to determine U's result will produce different bit-strings according to their probability: A single execution of a quantum algorithm is like a random experiment. Because of this, a quantum algorithm is typically performed many times to produce a probability distribution of results (see Figure 2 for an example) – and the most probable result is taken as "the" result of the quantum algorithm.



Fig. 2. Depicting a qubit and its manipulation (see [44] Fig. 2)

#### 2.3 Quantum Algorithms

As shown in Figure 3, the core of a quantum algorithm is a unitary transformation – which represents the proper logic of the algorithm. Its input register  $|r\rangle$  is prepared in a separate step (which turns out to be surprisingly complex [60, 73, 76]). Once the unitary transformation produced its output  $|s\rangle$ , a separate measurement step determines its result.



Fig. 3. Basis structure of a quantum algorithm (see [44] Fig. 3)

Optionally, some pre-processing or some post-processing is performed in a classical environment turning the overall algorithm into a hybrid one. Especially, many successful algorithms in a NISQ environment make use of classical processing to reduce the execution time on a quantum computer: the goal is to avoid decoherence and gate faults by spending only a short amount of time on a noisy quantum machine.

One example is a hybrid algorithm called Variational Quantum Eigensolver for determining eigenvalues [59]. This can be done by using a parameterized quantum algorithm computing and measuring expectation values, which are post-processed on a classical computer. The post-processing consists of a classical optimization step to compute new parameters to minimize the measured expectation values. The significance of this algorithm lies in the meaning of eigenvalues for solving many practical problems (see section 5.3).

Another example is the Quantum Approximate Optimization Algorithm [24] that is used to solve combinatorial optimization problems. It computes a state on a quantum

machine the expectation values of which relate to values of the cost function to be maximized. The state is computed based on a parameterized quantum algorithm, and these parameters are optimized by classical algorithms in a post-processing step as before. Since many machine learning algorithms require solving optimization problems, the importance of this algorithm is obvious too (see section 5.3).

An overview on several fundamental (non-hybrid) algorithms can be found in the work by Montanaro [51].

## 2.4 Quantum Software Stack

Programming a quantum computer is supported by a software stack the typical architecture of which is shown in Figure 4. LaRose describes incarnations of this stack by major vendors [39]. Also, section 3 discusses details of some implementations.



Fig. 4. Principle architecture of today's quantum software stack (see [44] Fig. 4)

The heart of the stack is a quantum assembler: it provides a textual rendering for key unitary transformations that are used to specify a quantum algorithm.

Since a quantum assembler is very low level, quantum programming languages are offered that host the elements of the quantum assembler in a format more familiar to traditional programmers – but still, the assembler flavor is predominant. In addition, functions to connect to quantum machines (a.k.a. quantum processing unit QPU) and simulators etc. are provided.

Quantum programming languages also come with libraries that provide implementations of often used quantum algorithms to be used as subroutines.

A compiler transforms a quantum assembler program into an executable that can be run on a certain QPU. Alternatively, the compiler can transform the quantum assembler into something executable by a simulator on a classical CPU.

#### 2.5 Sample Research Questions

The most fundamental question is about a proper engineering discipline for building (hybrid) quantum applications. For example: What development approach should be taken? How do quantum experts interact with software engineers? How are quantum applications tested and debugged?

## 3 Quantum as a Service

Since quantum algorithms promise to speed up known solutions of several hard problems in computer science, research in the field of software development for quantum computing has increased in recent years. In order to achieve speedup against classical algorithms, quantum algorithms exploit certain quantum-specific features such as superposition or entanglement [37]. The implementation of quantum algorithms is supported by the quantum software stack as shown in Figure 4. In this section, we give an overview of current tools for the development of quantum software. We further discuss deployment, different service models, and identify open research areas.

#### 3.1 Tooling

Several platforms implementing the introduced quantum computing stack have been released in recent years [39]. This includes platforms from quantum computer vendors, such as Qiskit [64] from IBM or Forest [62] from Rigetti, as well as platforms from third-party vendors such as ProjectQ [78] or XACC [48].



Fig. 5. Example of a quantum circuit (see [44] Fig. 5)

The quantum algorithms are described by so-called quantum circuits which are structured collections of quantum gates. These gates are unitary transformations on the quantum register (see section 2.3). Each platform provides a universal set of gates that can be used to implement any quantum algorithm. Figure 5 shows a simple example of such a circuit. It uses two qubits (each represented as a horizontal line), both of which are initialized as  $|0\rangle$ . A classical two-bit register c is used for the results of measurement and depicted as one single line. The Hadamard gate (H), which creates an equal superposition of the two basis states  $|0\rangle$  and  $|1\rangle$ , is applied to the qubit at quantum register position 0. Then, the Controlled Not gate (CNOT) is applied to the qubits at quantum register positions 0 and 1, whereby the former acts as control-bit and a NOT operation is applied to the second qubit iff the control qubit is  $|1\rangle$ . Finally, measurement gates are added to both qubits stating that these qubits will be measured and the resulting values will be stored in the classical bit register.

The different platforms support different quantum programming languages, which are embedded in classical host languages, such as PyQuil from Forest embedded in Python, or Qiskit embedded in Python, JavaScript, and Swift. The platforms provide libraries with methods for implementing a quantum circuit. The listing below shows a code snippet example of the creation and execution of the circuit from Figure 5. The

first line imports the library. Then, a circuit object is created to accumulate the gates in sequential order. Gate H is added to the circuit in line 4 and the CNOT gate is added to the circuit in line 5. Finally, measurement is added to the circuit in line 9. After the circuit is built, a concrete backend is chosen in line 11, which can be either a local simulator, a simulator in the cloud, or a QPU. The execution of the circuit is initiated in line 14. This execute method requires the circuit, the chosen backend, and the number of shots as input. As stated in section 2.2, a quantum algorithm is normally executed multiple times and the number of executions can be configured using the shots parameter.

```
1
   from SDK import lib
2
   # create circuit and add gates
3
   circuit = lib.Circuit()
4
   circuit.H(0)
5
   circuit.CNOT(0, 1)
6
   # ...
7
   # many more
8
   # ...
9
   Circuit.measure()
10
   # choose QPU
   Backend = lib.getBackend('...')
11
12
   # compile circuit and send to QPU
   result = lib.execute(circuit, backend, shots)
13
```

The circuit is then converted to quantum assembler language by the complier of the respective platform, e.g., to OpenQASM for QPUs of IBM, or Quil [77] for QPUs of Rigetti. In section 4.4 quantum compilers are introduced in more detail. The compiled code is sent to the selected backend. The execution itself normally is job-based, meaning that it will be stored in a queue before it gets eventually executed. The result, as mentioned before, is a probability distribution of all measured register states and must be interpreted afterwards.

Although the vendor-specific libraries are embedded in high-level programming languages, the implementation of quantum algorithms using the universal sets of gates requires in-depth quantum computing knowledge. Therefore, libraries sometimes already provide subroutines for common quantum algorithms, such as the Variational Quantum Eigensolver, or Quantum Approximate Optimization Algorithm. LaRose [39] compares different libraries with regards to their provided subroutines. However, these subroutines can often not be called without making assumptions about their concrete implementation and the used QPU.

Currently, most platforms are provided by the quantum computer vendors and are, thus, vendor-specific. However, there are also vendor-agnostic approaches, such as ProjectQ or XACC that both are extensible software platforms allowing to write vendoragnostic source code and run it on different QPUs. Section 4 gives more details on the hardware-agnostic processing of quantum algorithms.

#### 3.2 Deployment and Quantum Application as a Service

Several quantum computer vendors provide access to their quantum computers via the cloud. This cloud service model can be called Quantum Computing as a Service (QCaaS) [65]. Also cloud providers, such as Amazon, Microsoft, or 1Qbit, have taken QCaaS offerings to their portfolio. The combination of quantum and traditional computing infrastructure is essential for the realization of quantum applications. As already shown in Figure 3, a quantum computer is typically not used on its own but in combination with classical computers: the latter are still needed to store data, pre- and post-process data, handle user interaction etc. Therefore, the resulting architecture of a quantum application is hybrid consisting of both quantum and classical parts.

The deployment logic of the quantum part is currently included in the source code as shown in the code snippet from section 3.1. For running a quantum application (i) the respective platform has to be installed on a classical computer, (ii) the circuit must be implemented, (iii) the backend has to be selected, and (iv) the circuit must be executed. Therefore, we propose another service model that we call Quantum Application as a Service (QaaS), which is depicted in Figure 6. The QaaS offering wraps all application and deployment logic of a quantum application, including the quantum circuit as well as data pre- and post-processing, and provides an application programming interface (API) that can then be used for the integration with traditional applications, e.g., web applications or workflows.



**Fig. 6.** Quantum Application as a Service (QaaS) and Quantum Computing as a Service (QCaaS) (see [44] Fig. 6)

The traditional application passes input data to the API. However, this input data must be properly encoded in order to initialize the quantum register for the following computation [40]. This data encoding, the construction of an appropriate quantum circuit, its compilation, and the deployment is all handled by the service. For the execution of the circuit itself a QCaaS offering can be used. A hardware-agnostic processing of

quantum algorithms would also enable the flexible selection of different QCaaS as further discussed in section 4. The result of this execution is interpreted by the quantum application and finally returned to the traditional application.

This concept would enable to separate quantum applications from traditional applications, particularly with regard to their deployment. Furthermore, the integration of quantum computing features can be eased since QaaS enables to use common technologies of service-based architectures.

## 3.3 Sample Research Questions

To realize the proposed concept, the driving question is: How are hybrid quantum-classical applications deployed? In addition, the integration of quantum applications with traditional applications must be considered. This raises further questions. For example: What are the details of quantum algorithms, and especially their input and output formats? What are efficient encodings of input data? And for which parts of an application can a speedup be achieved?

# 4 Removing Hardware Dependencies

In this section, we motivate the need for removing the dependencies of quantum algorithms from quantum hardware and vendor-specific quantum programming languages. Afterwards, we present a method for the processing of hardware-independent quantum algorithms. Further, we sketch existing approaches to compile quantum algorithms to executables, optimize them, and show open research questions for selecting and distributing the quantum algorithms over suitable quantum and classical hardware.



Fig. 7. Processing of hardware-independent quantum algorithms (see [44] Fig. 7)

## 4.1 Problem

Due to the rapid development and improvement of quantum computers [54], it is important to keep implementations of quantum algorithms as hardware-independent and portable as possible, to enable the easy exchange of utilized quantum machines. Novel quantum algorithms are mostly specified and published in the abstract quantum circuit representation [80]. Therefore, to execute them, they must be implemented using the quantum programming language of a specific vendor (see section 3.1). However, the

quantum programming languages are not standardized and are usually only supported by a small subset or even only one quantum hardware vendor [39]. Therefore, the implementation of a quantum algorithm utilizing a specific quantum programming language can lead to a vendor lock-in. To circumvent this problem, a standardized, machine-readable, and vendor-agnostic representation for quantum circuits is required, which can be automatically translated into the representations of the different vendorspecific quantum programming languages (see section 2.4).

Furthermore, after specifying a quantum algorithm using a certain quantum programming language, the utilized qubits and gates must be automatically mapped to qubits, gates, and measurements that are provided by the quantum machine to keep them independent of different quantum machines of a specific vendor [13].

## 4.2 Hardware-Independent Processing

In this section, we present a method for the processing of hardware-independent quantum algorithms, which is based on the works of Häner et al. [29] and McCaskey et al. [48]. First, the required steps are presented and afterwards the following sections introduce available research works that can be integrated into the approach and provide an overview of open research questions for the different steps.

The required processing steps for hardware-independent quantum algorithms are sketched in Figure 7. The inputs and outputs of the different steps are depicted by the arrows connecting them. First, the quantum algorithm is defined utilizing a vendor-agnostic quantum programming language, which should be standardized and comprise all relevant parts of quantum algorithms [48]. Then, a hardware-independent optimization can be performed (see section 4.5), which, e.g., deletes unnecessary qubits or gates [29].

Based on the optimized quantum algorithm, suitable quantum hardware is selected in the next step. For this, important properties characterizing the quantum algorithm, such as the required number of qubits or the utilized gate set, are retrieved [79]. Due to the limited quantum hardware in the NISQ era [61], this information is important and can be used to select a quantum computer that can successfully execute the quantum algorithm. Furthermore, this selection can be based on different metrics, such as the error-probability, the occurring costs, or the set of vendors that are trusted by the by the user [48].

After the selection of the quantum hardware to execute an algorithm, the algorithm must be translated from the vendor-agnostic quantum programming language to the quantum assembler of a vendor that supports the execution on the selected quantum hardware [48]. Next, it can be compiled to an executable for the selected quantum hardware. For this, the available vendors usually provide suitable compilers (see section 4.4) [39]. During the compilation process, hardware-dependent optimizations are performed. Finally, the executable can be deployed and executed on the selected quantum machine (see section 3.2).

#### 4.3 NISQ Analyzer

The NISQ Analyzer [69] is a component which analyzes quantum algorithms and extracts the important details, such as the number of required qubits or the utilized gate set [79]. Therefore, the quantum algorithm specified in the hardware-independent quantum programming language can be used as an input for the NISQ Analyzer. However, the analysis of quantum algorithms and the precise estimation of resource requirements are difficult problems [71]. For example, the required gates for the initial data encoding [40] or the overhead due to required error correction codes [38] must be considered. Additionally, the resource requirements for oracle implementations are often ignored but lead to a large overhead that should be noted [70]. Thus, tooling support is required that extracts all relevant characteristics of quantum algorithms and provides them to the other components, such as the quantum compiler.

### 4.4 Quantum Compiler

The quantum compiler is in charge of performing the mapping from the quantum assembler representing a quantum algorithm to an executable for a concrete quantum computer [13, 30]. The mapping of gates and measurements that are physically implemented by a quantum computer can be performed directly. However, gates and measurements that are not physically available have to be mapped to a "subroutine" consisting of physical gates and measurements [30]. For example, if a measurement using a certain basis is not implemented, the quantum state must be transferred into a basis for which a measurement is provided by the quantum hardware and the measurement must be done in this basis. The utilized subroutines strongly influence the execution time and error probability of the calculation, as they add additional gates and measurements [35, 78]. Hence, suited metrics and algorithms to select the required subroutines are important to reduce the overhead of the mapping (see section 4.5). Additionally, the qubits must be mapped to available physical qubits, which influences the quantum algorithm execution as well, due to different characteristics of the qubits, such as decoherence time or connectivity [87]. However, the available quantum compilers are mostly vendor-specific [39], and therefore, compile the quantum algorithm implementations defined in the quantum assembler of a certain vendor to the executable for concrete quantum hardware that is provided by this vendor. Other quantum compilers define their own quantum assembler language to specify quantum algorithms and map them to executables for a certain quantum computer as well [36]. Thus, the dependency on the vendor- or compiler-specific quantum assembler language cannot be removed by these kinds of quantum compilers. Hence, quantum compilers must be integrated into the approach for processing hardware-independent quantum algorithms (see Figure 7).

## 4.5 Optimization of Quantum Algorithms

Quantum algorithms can be optimized in two ways: (i) hardware-independent or (ii) hardware-dependent [29]. For the hardware-independent optimization, general optimizations at the quantum circuit level are performed, according to a cost function, such as

the circuit size or the circuit depth [80]. In contrast, hardware-dependent optimization takes hardware-specific characteristics, such as the available gate set of the target quantum computer or the decoherence time of different qubits, into account [34]. Hence, this optimization is often combined with the compilation to an executable for a certain quantum computer.

In the following, we sketch some existing works regarding the optimization of quantum algorithms. Heyfron and Campbell [30] propose a quantum compiler that reduces the number of T gates, while using the Clifford + T gate set. They show that the cost of the T gate is much higher than for the other Clifford gates, and therefore, they improve the circuit costs by decreasing the T count. Itoko et al. [34] present an approach to improve the hardware-dependent mapping from the utilized qubits and gates in the quantum algorithm to the provided qubits and gates of the quantum computer during the compilation process. Maslov et al. [47] propose an approach that is based on templates to reduce the circuit depth, which means the number of gates that are executed in sequence on the qubits. A template is a subroutine that can be used to replace functionally equivalent circuit parts by more efficient ones in terms of different metrics like cost or error probability. Hence, they introduce a method to detect and replace suitable circuit parts with templates.

## 4.6 Sample Research Questions

For the definition and processing of hardware-independent quantum algorithms and the selection of suitable quantum hardware, different research questions must be solved, some of which are presented in the following.

The definition of an abstract hardware-independent quantum programming language is important to remove the hardware dependencies of quantum algorithms. Therefore, sample research questions are: What elements are required to define quantum algorithms? How should suited modeling tooling support look like? What subroutines are important and should be provided as libraries?

To automatically select the best available quantum hardware for a quantum algorithm, suited tooling support must be developed. Hence, open research questions are: What characteristics of quantum algorithms are important for the hardware selection? How can these characteristics be retrieved automatically? What are suited metrics and algorithms for the hardware selection? What are the interesting optimization goals?

The hardware-dependent and -independent optimization of quantum algorithms are especially important in the NISQ era. Therefore, interesting research questions are: What are new or improved optimization algorithms? What data about quantum hardware is relevant for the optimization and how can it be obtained?

By comparing the performance of different quantum compilers, the compiler with the best optimization result or best execution time can be selected. Hence, sample research questions are: What are suited benchmarks for the comparison of quantum compilers? How can the optimality of the compiled executable be verified with respect to different optimization goals, like the number of required gates or the number of fault paths?

# 5 Quantum Humanities: A Use Case from Quantum Machine Learning

Determining how quantum computing can solve problems in machine learning is an active and fast-growing field called quantum machine learning [75]. In this section we give a use case from the digital humanities [10] that shows how quantum machine learning can be applied.

## 5.1 Quantum Humanities

As the following use case will stress, there are promising application areas for quantum computing not only in industry or natural science but also in the humanities. We coined the term *quantum humanities* for using quantum computing to solve problems in this domain [9]. It aims at exploiting the potentials offered by quantum computers in the humanities, raise research questions, and describe problems that may benefit from applying quantum computers. Figure 8 gives an overview of the process and algorithms (see section 5.3-5.5) used to analyze the data from our project called MUSE.



Fig. 8. MUSE data analysis (see [44] Fig. 9)

## 5.2 MUSE

MUSE [7, 52] aims at identifying costume patterns in films. Costume patterns are abstract solutions of how to communicate certain stereotypes or character traits by e.g. the use of specific clothes, materials, colors, shapes, or ways of wearing. To determine the conventions that have been developed to communicate for example a sheriff or an outlaw, MUSE developed a method and a corresponding implementation to support the method to capture and analyze costumes occurring in films.

The method consists of five main steps: (1) defining the domain by an ontology, (2) identifying – based on strict criteria – the films having most impact within the domain, (3) capturing all detailed information about costumes in films in the MUSE repository,

(4) analyzing this information to determine costumes that achieve a similar effect in communicating with the recipient, and (5) abstracting these similarities to costume patterns [4, 7]. This method has been proven to be generic by applying it in our parallel project MUSE4Music [6].

**Ontology.** To structure costume parameters that have a potential effect on the recipient of a film a detailed ontology was developed [4, 5]. As a basis for this ontology, several taxonomies have been devised, which structure sub-domains such as kinds of clothing, materials, functions or conditions. An example of such a taxonomy can be seen in Figure 9: The taxonomy of colors structures a relevant subset of colors in hierarchical form (see section 5.4 on how these taxonomies contribute to determine of similarity measures between costumes).



Fig. 9. Taxonomy of colors (based on [5])

The ontology brings these taxonomies together and adds relations (e.g. worn above, tucked inside, wrapped around etc.) on how base elements (e.g. trousers, shirts, boots etc.) are combined into an overall outfit. The 3151 nodes of the ontology induce the schema of the MUSE repository. The repository facilitates the structured capturing of all relevant information about the films, their characters and their costumes.

**Data Set.** The MUSE data set currently (July 2020) contains more than 4,900 costumes out of 58 films, consisting of more than 27,000 base elements, 60,000 primitives (e.g. collar, sleeves, buttons, etc.), 151,000 color selections and 171,000 material selections.

In the sense of the open data initiative, this data set is freely available to be used and analyzed [53]. It provides a well-structured and labelled data set that allows several analysis techniques to be applied.

**Data Analysis.** As a first approach to analyze the data to identify those significant elements a costume designer uses to achieve a certain effect, a two-step analysis process was introduced [19]. The first step applies data mining techniques – mainly association rule mining – to determine hypotheses about which elements are used to communicate a certain stereotype, for example. The second step aims at refining and verifying such hypotheses by using online analytical processing (OLAP) techniques [23] to identify indicators for costume patterns.

To improve the process of building hypotheses that hint to potential costume patterns we are currently extending the analysis of the MUSE data by various techniques from machine learning. Each costume is described by a huge number of properties. Simply mapping each property of a costume to a feature, the resulting feature space would be of huge dimension. Therefore, feature extraction, namely principle component analysis (PCA), multidimensional scaling (MDS) and autoencoders, are applied to reduce the dimension of the feature space without losing important information. To group those costumes together, that achieve the same effect, different cluster algorithms, such as maximum cut, k-means etc. are applied and evaluated. As there are new costumes frequently stored into the database the usage of classification algorithms is investigated to enable classifying these costumes.

Currently, this approach is implemented on a classical computer with classical machine learning algorithms. But since quantum computing can contribute to solve several problems in machine learning – as shown in the following section – it is promising to improve the approach accordingly [8].

### 5.3 Potential Improvements

Several machine learning algorithms require the computation of eigenvalues or apply kernel functions: these algorithms should benefit from improvements in the quantum domain. Many machine learning algorithms are based on optimization, i.e. improvements in this area like Quantum Approximate Optimization Algorithm QAOA should imply improvements of those machine learning algorithms.

Whether or not such improvements materialize is discussed in several papers that compare sample classical and quantum machine learning algorithms, e.g. [11, 14, 28].

**Data Preparation.** The data captured in MUSE are categorical data mostly. Since most machine learning algorithms assume numerical data, such categorical data must be transformed accordingly: this is a complex problem.

For example, the different colors of pieces of clothes could be assigned to integer numbers. But the resulting integers have no metrical meaning as required by several machine learning algorithms. Instead of this, we exploited the taxonomy that structures all of our categorical data by applying the Wu and Palmer metric [85] to derive distances between categorial data (for more details see section 5.4).

As described above, costumes have a large number of features, thus, this number must be reduced to become tractable. We experiment with feature extraction based on restricted Boltzmann machines [31, 32] as well as with principal component analysis,

embeddings and autoencoders (see section 5.5). Feature selection based on deep Boltzmann machines [81] may also be used.

**Eigenvalues.** Principal component analysis strives towards combining several features into a single feature with high variance, thus, reducing the number of features. For example, in Figure 10 the data set shown has high variance in the A axis, but low variance in the B axis, i.e. A is a principal component. Consequently, the X and Y features of the data points are used to compute A values as a new feature, reducing the two features X and Y into a single feature A.



Fig. 10. Principal component of a data set (see [44] Fig. 8)

The heart of this analysis is the calculation of the half axes and their lengths of the ellipse "best" surrounding the data set. This is done by determining the eigenvalues of the matrix representing the ellipse. Computing eigenvalues can be done on a quantum computer much faster than classically by means of quantum phase estimation and variational quantum eigensolvers. Thus, Quantum principal component analysis [45] is an algorithm we will use in our use case.

**Quantum Boltzmann Machines.** Zhang, Li and Zhu [86] provided a quantum algorithm of a quantum restricted Boltzmann machine. In a use case, it has shown performance superior to a classical restricted Boltzmann machine.

Similarly, Amin et al. [2] described an approach for both, quantum Boltzmann machines as well as quantum restricted Boltzmann machines. They report that the quantum restricted Boltzmann machine outperforms the classical restricted Boltzmann machine for small size examples.

Thus, quantum Boltzmann machines are candidates for our use case, especially because they can be exploited in clustering and classification tasks.

**Clustering.** Several quantum clustering algorithms and their improvements over classical algorithms are presented by Aimeur, Brassard and Gambs [1]. Since clustering can be achieved by solving Maximum Cut problems, some attention has been paid to solve MaxCut on quantum computers.

For example, Crooks [17] as well as Zhou et al. [88] use QAOA to solve MaxCut problems on NISQ machines. A similar implementation on a Rigetti quantum computer has been described by Otterbach et al. [58]. Thus, quantum clustering is promising.

**Classification.** Support vector machines (SVM) are established classifiers. Rebentrost, Mohseni and Lloyd [66] introduce quantum support vector machines and show an exponential speedup in many situations.

Schuld, Sinayskiy and Petruccione [74] present a quantum version of the k-nearest neighbor algorithm, and an implementation of a classifier on IBM Quantum Experience [72]. A hybrid classifier has been introduced by Schuld and Killoran [73].

The use of kernels in machine learning is well-established [33] and kernels are used in case non-linear separable data must be classified. A hybrid classifier that makes use of kernels is given by Schuld and Killoran [73]. Ghobadi, Oberoi and Zahedinejhad [25] describe classically intractable kernels for use even on NISQ machines. Thus, quantum classifiers are promising.

#### 5.4 Categorical Data

Before the potential improvements of quantum machine learning outlined above can be applied to our use case, we need to address the problem of categorical data, since most of the MUSE data is categorical and the relevant algorithms use numerical data. As our use case can be seen as paradigmatic for further use cases from quantum humanities, addressing the problem of categorical data will be discussed in more detail in this section.

Categorical data is data with a finite set of values, e.g. enumerations of strings. If it has a canonical order it is called ordinal data, otherwise it is called nominal data. Ordinal data can be compared with "greater than" or "less than", for example, while nominal data only support comparisons with "equal" and "not equal". Calculations on categorical data, even if it is numerical, have in general no meaning. For example, what is the maximum value of a set of colors? What is the sum of ZIP codes? What is the average of the two dress sizes M and XXL?

As stated above, many algorithms in machine learning expect metrical data as input, i.e. numerical data that support calculations and comparisons in a meaningful manner. For example, the average of a set of salaries can be computed, its mean value etc. and the result has a well-understood semantics. Such calculations are often performed to compute the distance of data points in a feature space, e.g. to determine clusters of data points or whether a new data point belongs to a known cluster.

Thus, to benefit from corresponding algorithms, categorical data has to be transformed into metrical data. A very basic transformation may assign a natural number to a categorical data item like "red  $\mapsto$  1" and "blue  $\mapsto$  5". But algorithms would then infer that "red < blue" which is not meaningful. There are transformations that avoid such problems (like one-hot-encoding) by assigning each different values of a categorical parameter to a new dimension in a feature space. As a consequence, if a parameter (like ZIP code) has many different values or if many categorical parameters must be processed, such encodings create feature spaces with very high dimensions — which demand huge processing power.

Embeddings try to solve this problem by mapping a high-dimensional feature space F into a significantly lower-dimensional space. The latter space is the normed vector space ( $\mathbb{R}^n$ ,  $\|\cdot\|$ ), where n  $\ll$  dim F and  $\|\cdot\|$  is any proper norm on  $\mathbb{R}^n$ . The embedding

is a map  $\varphi : F \to \mathbb{R}^n$  such that the distance ||x - y|| between two points  $x = \varphi(a)$  and  $y = \varphi(b)$  approximately corresponds to the similarity of the original data points  $a, b \in F$ .

#### Measuring Similarity of Parameters with Tree-Structured Domains.

<u>Definition</u>: For a finite set M a *similarity measure on* M is a map  $s: M \times M \rightarrow \mathbb{R}$  with the following properties:

- i. s(i,j) = s(j,i) (symmetry)
- ii.  $s(i,i) \ge s(j,i)$  (maximal similarity)
- iii.  $s(i,j) \ge 0$  (positivity)
- iv. s(i,i) = 1 (self-identity)

where i, j are arbitrary elements in M. Often,  $s_{ij}$  is written instead of s(i,j).

For example, when analyzing text corpora, the cosine similarity is typically used that assigns the cosine between two vectors each of which represents a document [46] as similarity of two documents. In case the domain of a categorical parameter is structured by means of a taxonomy (or more general, by means of a tree) the Wu-Palmer similarity measure [85] can be used. To define this measure, we remind three definitions.

<u>Definition</u>: A finite graph G = (N, E) is called a *tree* iff the following holds:

- i. G is directed
- ii. card E = card N 1
- iii.  $\exists ! w \in N : d_{in}(w) = 0$

The node w is called *root* of G ( $d_{in}(n)$  is the number of incoming edges of node n, i.e. the root w has no incoming edges).  $\Box$ 

The Wu-Palmer similarity measure is based on length of paths in the tree structuring the domain of a categorical parameter.

<u>Definition</u>: Let x,y ∈ N be two nodes of a tree G. A *path*  $p_{x,y}$  *from x to y* is a set of edges {(x<sub>0</sub>, x<sub>1</sub>), (x<sub>1</sub>, x<sub>2</sub>)...,(x<sub>n-1</sub>, x<sub>n</sub>)} ⊆ E with x=x<sub>0</sub> and x<sub>n</sub>=y. A path is often specified as  $p_{x,y} : x=x_0 \rightarrow x_1 \rightarrow ... \rightarrow x_{n-1} \rightarrow x_n=y$  and the number of edges of a path is called the *length* of the path L(p<sub>x,y</sub>). □



Fig. 11. A path between the nodes x and y in a tree

Figure 11 depicts a tree and highlights the path  $p_{x,y}$  in the tree between nodes x and y. The length  $L(p_{x,y})$  of this path is 3:  $L(p_{x,y})=3$ . Note, that the direction of edges of a tree are often not explicitly drawn because their direction is assumed to be from top to bottom.

<u>Definition</u>: Let G = (N, E) be a tree, w its root,  $x, y \in N$ , and let  $p_{x,w}: x=x_0 \rightarrow x_1 \rightarrow \dots \rightarrow x_{n-1} \rightarrow x_n=w$  and  $p_{y,w}: y=y_0 \rightarrow y_1 \rightarrow \dots \rightarrow y_{m-1} \rightarrow y_m=y$  be the paths from x to w and from y to w, respectively. The uniquely defined  $v \in N$  with the properties

i.  $v \in \{x_0, x_1, \dots, x_{n-1}, x_n\} \cap \{y_0, y_1, \dots, y_{m-1}, y_m\}$ ii.  $\forall (v, v') \in E: v' \notin \{x_0, \dots, x_n\} \cap \{y_0, \dots, y_m\}$ is called the *lowest common ancestor* of x and y.  $\Box$ 



Fig. 12. The lowest common ancestor of two nodes in a tree, and paths used to determine the Wu-Palmer similarity measure

The lowest common ancestor v of the two nodes x and y are depicted in Figure 12. The two paths  $p_{x,w}$  and  $p_{y,w}$  from x and y to the root are shown in the left side of the figure, and v is the lowest common ancestor of x and y. In the right hand side of the figure, the paths  $p_{x,w}$  and  $p_{y,w}$  are split into paths  $p_{x,v}$  and  $p_{y,v}$  from x and y to their lowest common ancestor v, and another path  $p_{v,w}$  from v to w. These paths are used next to finally define the Wu-Palmer similarity measure [85] as follows:

<u>Definition</u>: Let G = (N, E) be a tree, w its root,  $x,y \in N$ , and v the lowest common ancestor of x and y. Then,

$$\omega(x, y) := \frac{2 \cdot L(p_{v,w})}{L(p_{x,v}) + L(p_{y,v}) + 2 \cdot L(p_{v,w})}$$

is called the *Wu-Palmer similarity* of x and y.  $\Box$ 

Note, that the path  $p_{x,w}$  can be written as  $p_{x,w} : p_{x,v} \rightarrow p_{v,w}$ , i.e. as the composite path from x to v followed by the path from v to w (see Figure 12), and similar for the path  $p_{y,w}$ . With this it is obvious that the Wu-Palmer similarity can be computed as

$$\omega(x, y) = \frac{2 \cdot L(p_{v,w})}{L(p_{x,w}) + L(p_{y,w})}$$

The following can be seen directly from the definition of a similarity measure:

<u>Proposition</u>: The Wu-Palmer similarity  $\omega$  is a similarity measure on the node set of a tree G.

Thus, categorical parameters the domains of which are structured by a tree (e.g. by a taxonomy) can be measured with respect to their similarity. As an example, the similarity of x and y in Figure 12 is

$$\omega(x, y) = 2 \cdot \frac{L(p_{v,w})}{L(p_{x,v}) + L(p_{y,v}) + 2 \cdot L(p_{v,w})} = 2 \cdot \frac{2}{3 + 3 + 2 \cdot 2} = \frac{2}{5}$$

Similarity of Set-Valued Categorical Parameters. A categorical parameter may be set-valued. E.g. a shares portfolio may contain shares from different companies. Any similarity measure s on a given set N can be extended to a similarity measure on the powerset of N,  $\wp(N)$ . Thus, categorical parameters with set-valued domains can be measured in terms of their similarity if the (single) domain the members of the set come from are equipped with a similarity measure.

<u>Definition</u>: Let G = (N, E) be a tree, A, B  $\subseteq$  N, and s be a similarity measure on N. Then,

$$\sigma(A,B) := \frac{1}{2} \left( \frac{1}{cardA} \cdot \sum_{a \in A} \max_{b \in B} (a,b) + \frac{1}{cardB} \cdot \sum_{b \in B} \max_{a \in A} (a,b) \right)$$

defines a map  $\sigma: (\wp(N) \setminus \emptyset) \times (\wp(N) \setminus \emptyset) \to \mathbb{R}$ .  $\Box$ 

The following follows from the fact, that s is a similarity measure: Proposition:  $\sigma$  is a similarity measure.

Figure 13 shows two sets  $A \subseteq N$  and  $B \subseteq N$ . The similarity values of pair from  $A \times B$  are given as annotations of the dashed lines indicating the proper pairs, e.g. s(a,y)=0.3. The maximum value of the pairs with the same first component are noted close to the first component. Based on this the similarity  $\sigma(A,B)$  of the two sets A and B is computed as

$$\sigma(A,B) = \frac{1}{2} \left( \frac{1}{3} \cdot (0.7 + 0.9 + 0.8) + \frac{1}{2} \cdot (0.8 + 0.9) \right) = \frac{1}{2} (0.8 + 0.85) \approx 0.825$$



Fig. 13. Comparing set-valued categorical parameters

Especially, if a domain N is structured by a tree, the Wu-Palmer similarity measure induces a similarity measure on  $\mathcal{P}(N)$ . This can be extended to tuples of set-valued parameters:

<u>Definition</u>: Let N<sub>i</sub> be sets with similarity measure s<sub>i</sub>, and let  $\sigma_i$  be the similarity measure on  $\mathcal{P}(N_i)$  induced by s<sub>i</sub>. Then,

$$\mu((A_1, \dots, A_m), (B_1, \dots, B_m)) := \frac{1}{m} \sum_{i=1}^m \sigma_i(A_i, B_i)$$

defines a map

$$\mu: \prod_{i=1}^{m} (\wp(N_i) \setminus \emptyset) \times \prod_{i=1}^{m} (\wp(N_i) \setminus \emptyset) \to \mathbb{R}. \ \square$$

Directly from the definitions, the following is implied: <u>Proposition</u>:  $\mu$  is a similarity measure.

**Costume Similarity Matrix.** The data from the MUSE database are mostly categorical. Based on the taxonomies of the costume parameters that are used to capture all relevant information about a certain costume (see section 5.2) the similarity of the costumes and their base elements can be determined as described above.

Often, algorithms are more conveniently formulated based on the distance of data than based on similarity. For this purpose, the notion of distance must be defined.

### 5.5 Distance Measures and Feature Extraction

<u>Definition</u>: Let M be a finite set. A *distance measure* (a.k.a. *dissimilarity measure*) on M is a map  $t: M \times M \to \mathbb{R}$  with the following properties:

- i. t(i,j) = t(j,i) (symmetry)
- ii.  $t(i,j) \ge 0$  (positivity)
- iii.  $t(i,j) = 0 \Leftrightarrow i = j$  (definiteness)

where i, j are arbitrary elements in M. Often,  $t_{ij}$  is written instead of t(i,j). In case t in addition has the property

iv.  $t(i,j) \le t(i,k) + t(k,j)$  (triangle inequality) for arbitrary elements i, j, k in M, t is called a *metric* on M.  $\Box$ 

A categorical parameter the domain of which is structured by a tree has a distance measure defined that can be derived from the Wu-Palmer similarity measure. This is because similarity measures induce distance measures and vice versa:

Proposition: If s is a similarity measure on M, then

$$t(i,j) = \sqrt{s(i,i) + s(j,j) - 2s(i,j)}$$

is a distance measure on M. If t is a distance measure on M, then

$$s(i,j) = \frac{1}{1+t(i,j)}$$

is a similarity measure on M.

For example, in Figure 12 the similarity of x and y is  $\omega(x,y) = 2/5$ . Thus, their distance based on the induced distance measure from the proposition is  $t(x,y) = \sqrt{6/5}$ . Note, that there are different ways to turn a similarity measure into a distance measure and vice versa. The proposition only gives two straightforward ways.



Fig. 14. Distance matrix for color and genre values of 21 random chosen costumes

Figure 14 depicts a distance matrix for the color and genre values of 21 random chosen costumes from the database using the Wu-Palmer similarity measure. As most of the

costumes have more than one color or genre value the similarity of set-valued categorical parameters is used. The color code in Figure 14 corresponds to the degree of distance of two costumes and is described on the right side of the figure: dark blue corresponds to a distance value of 0, i.e. the two elements are equal, and dark red, with a distance value of 1, indicates maximum dissimilarity. Based on distance measures the concept of an embedding can now be precisely defined.

<u>Definition</u>: Let M be a finite set and t be a distance measure on M. A map  $\varphi: (M, t) \rightarrow (\mathbb{R}^n, d)$  with the property  $d(\varphi(i), \varphi(j)) \approx t(i, j)$  for all  $i, j \in M$  is called an *embedding*. Here, d is any appropriate metric on  $\mathbb{R}^n$ .  $\Box$ 

Several algorithms are known to compute embeddings. Especially when data sets should be visualized, embeddings in two- or three-dimensional space (i.e. n = 2 or n = 3) are desirable. For this purpose, multidimensional scaling (MDS) algorithms [16] are used to compute embeddings. Figure 15 depicts the result of an embedding. If  $M \subseteq \mathbb{R}^m$  and  $m \gg n$ , an embedding is a feature extraction technique for dimensionality reduction to make processing data in machine learning tractable.



Fig. 15. An embedding of a set with distance metric into  $\mathbb{R}^2$  with Euclidian metric

As an example of MDS applied to the presented use case, Figure 16 (left side) depicts the embedding of the color and genre values of 40 randomly selected costumes and their distances into the two-dimensional space with the Euclidean metric.



Fig. 16. Multidimensional scaling (left side) and clustering (right side)

This embedding allows to run different cluster algorithms to determine those costumes that have the same effect. Figure 16 (right side) gives an example on how such costume clusters may look like. The algorithm used is the OPTICS (Ordering Points to Identify the Clustering Structure) clustering algorithm that creates an augmented ordering of the neighboring points representing its density-based clustering structure [3].

As can be seen in this example, not all costumes (depicted as numbers) can be classified as part of the identified clusters (colored circles). By taking a closer look at the individual clustered costumes they indicate candidates of costume patterns, e.g. the cluster in the lower left corner (red circles) reveals that in the Western genre one of the dominant colors of the costumes is black. The candidate clusters require verification by a broader data set.

**Classical and Hybrid Autoencoder.** Besides PCA or embeddings, there are several other algorithms to perform feature extraction, e.g. autoencoders. In the following, we show how an autoencoder is used in a quantum-inspired environment to analyze the MUSE data.

An autoencoder is a neural network performing the task of reconstructing the original input data as output data. The strength of this approach is to reduce the dimensions of a feature space by learning a representation of the input training data set. Romero, Olson and Aspuru-Guzik [67] introduce a quantum autoencoder for efficient compression of quantum data. Based on this, we have implemented a quantum-classical hybrid autoencoder in Qiskit [64] and TensorFlow Quantum (TFQ) [83]. The quantum autoencoder performs five main steps: (i) prepare the input data, (ii) learn how to preserve the quantum information of the input data with a reduced number of qubits (corresponding to the dimensions to the feature space) by an encoder, (iii) reset the qubits that are supposed to be compressed to |0), (iv) reconstruct the quantum states of the input data as output by a decoder, and (v) measure the quantum states.



Fig. 17. Training sup-part of a quantum autoencoder circuit

Figure 17 depicts the sub-part of the implemented circuit run on Qiskit performing the training of the neural net. All five qubits are initialized to  $|0\rangle$ . The first three qubits  $q0_i$  represent three dimensions. The second two qubits  $q1_0$ ,  $q2_0$  are ancillae qubits that are used to perform a SWAP test. The R<sub>x</sub> gates represent rotations around the x-axis to achieve an angle encoding of the input data [41, 84]. The other R gates are rotations, which are parameterized by three values. Those values are initially randomly chosen and then modified via gradient decent in each iteration. Some of these rotations (pink) are controlled rotations. The final SWAP test determines how much the two quantum states  $q0_2$  and  $q1_0$  differ.

An autoencoder can be trained by minimizing the reconstruction error, i.e. the difference between the original input and the reconstructed output. In our use case we compared the reconstruction error of classical and hybrid autoencoders. During the first step a classical autoencoder was used to reduce all color values of the MUSE dataset to a three-dimensional feature space, while in the second step (i) classical autoencoders and (ii) quantum autoencoders were used to further reduce the dimension to two.

In Figure 18, the reconstruction error is visualized on the y-axis, and the performed number of iterations is depicted on the x-axis (average of 10 passes). The following autoencoders have been used: in our classical implementation we used (1) PyTorch [63] (orange line) and (2) TensorFlow [82] (green line); our hybrid implementation used (3) Qiskit and PyTorch (blue line) and (4) TFQ (yellow line). As several tests indicated that already the required SWAP test is highly error-prone due to the decoherence of the qubits and the fidelity of the gates on the quantum device, we have chosen a simulator to perform our hybrid encoder implemented with Qiskit and PyTorch.



Fig. 18. Reconstruction errors (y-axis) and the number of iterations (x-axis) of classical and hybrid autoencoders

The results depicted by Figure 18 show that the reconstruction error of the hybrid autoencoders are both smaller compared to the classical autoencoders (error rate per autoencoder: (1) 0.0698, (2) 0.0696, (3) 0.0681, (4) 0.0665) and require much fewer iterations than the classic autoencoders. However, the downside of our hybrid autoencoders, especially the autoencoder implemented in Qiskit and PyTorch, is their longer training time compared to the classical autoencoders (time in seconds per autoencoder: (1) 1,8, (2) 3,7, (3) 2660, (4) 28,9). Nevertheless, the hybrid implementations are preferable because their results are quantum states, i.e. they can be immediately processed by further quantum algorithms, e.g. quantum cluster algorithms, benefitting from the potential improvements by using the quantum computer (see section 5.3) accessible via a cloud (see section 3).

#### 5.6 Sample Research Questions

The most essential and fundamental question for quantum humanities is to evaluate which existing and new problems from the humanities can be addressed by quantum computers. Especially, which problems are best solved by classical, hybrid, or quantum algorithms? Besides speedup, which algorithms result in higher precision? Which language allows to communicate between many disciplines (e.g. mathematics, physics, computer science, and the different areas from the humanities)? Are there completely new questions from the humanities that are only addressable based on a quantum computer?

# 6 Collaborative Quantum Application Platform

Driven by the continuous improvement of quantum hardware, specialists in various fields have developed new quantum algorithms and applications in recent years. The use of these quantum applications requires in-depth knowledge of theory and practice, which is often lacking in small and medium-sized companies. A major challenge today is to facilitate the transfer of knowledge between research and practice to identify and fully exploit the potential of new emerging technologies. To prepare a body of knowledge for quantum computing reasonably and make it usable for different stakeholders, a collaborative platform where all participants come together is essential [43]. For this purpose, the quantum application platform must cover the entire process from the development of quantum algorithms to their implementation and execution. The diversity of stakeholders and their different objectives lead to a variety of requirements for such a quantum platform.

Building upon the stakeholders identified by Leymann, Barzen and Falkenthal [43], we firstly identify key entities, which serve as an anchor for the knowledge on a quantum platform, secondly identify essential requirements for their expedient implementation and, finally, show a general extendable architecture for a collaborative quantum software platform.

### 6.1 Key Entities

To foster a clear structuring of the knowledge created on a quantum software platform the following key entities can be used. They allow different experts to hook into the platform and enable to share and contribute knowledge.

**Quantum Algorithm.** As mentioned before, quantum algorithms are developed and specified typically by experts with in-depth quantum physics background. Thus, for a quantum software platform it is essential to capture quantum algorithms as artefacts. Besides generally sharing them, further valuable information can be attached to quantum algorithms, such as discussions among experts regarding resource consumption of an algorithm, its speedup against classical algorithms, or its applicability to NISQ computers.

Algorithm Implementation. Besides the representation of quantum algorithms in their conceptual form, i.e., as mathematical formulas or abstract circuits, the heterogeneous field of quantum hardware demands to capture vendor- and even hardware-specific implementations of quantum algorithms. This is because, implementations for a particular quantum computer offering of a vendor requires the use of a vendor-specific SDK. Thus, implementations of an algorithm for quantum computers offered by different vendors ends up in different code or even the usage of completely different quantum programming languages. Thus, enabling sharing of different algorithm implementations on a quantum software platform stimulates knowledge transfer and reduces ramp-up especially for unexperienced users.

**Data Transformator.** Since quantum algorithms rely on the manipulation of quantum states they do not operate directly on data as represented in classical software. Instead, the data to be processed must be encoded in such a way that they can be prepared into a quantum register. Different problem classes such as clustering or classification have specific requirements for the data to be processed. It can be of great benefit to identify general transformation and coding strategies for relevant problem classes. Such strategies can then be represented and discussed on the platform as data transformations.

**Hybrid Quantum Application.** Since only the quantum parts of an algorithm are executed on a quantum computer, they must be delivered together with classical software parts that run on classical computers. To exploit the full potential of quantum algorithms, they often have to be properly integrated into an already running system landscape, which includes proper data preparation and transformation. This is why solutions that are rolled out in practice are typically hybrid quantum applications (see section 3.2). Therefore, knowledge transfer about applicable software solutions for particular use cases at hand is bound to hybrid quantum applications.

**Quantum Pattern.** Software patterns are widely used to capture proven solution principles for recurring problems in many fields in computer science. Thus, quantum patterns seem to be a promising approach to also capture proven solutions regarding the design of quantum algorithms, their implementation and integration in existing systems. First patterns for developing quantum algorithms have already been published [40] and an entire pattern language on architecting quantum applications can be an object of investigation on the intended platform.

#### 6.2 Requirements

The essential challenge to create and provide a reasonable body of knowledge on quantum algorithms and applications involves the collaboration among several stakeholders. In contrast to traditional software engineering, quantum algorithms are typically not specified by computer scientist rather than by quantum physicists. Furthermore, to understand and implement those algorithms a different mindset is required because the key buildings blocks of algorithms are no longer loops, conditions, or procedure calls but quantum states and their manipulation via unitary operators.

By involving all participants identified by Leymann, Barzen and Falkenthal [43] in the platform, added value can be created, both for experienced quantum specialists and inexperienced customers. For this, the following listed requirements must be met.

**Knowledge Access.** Often only certain specialists and scientists have the required expertise for developing quantum algorithms and their implementation. To identify and exploit the use cases of quantum computing in practice, companies must be empowered to gather knowledge and to exchange with experts (developer, service provider, consultants, and so on) [50]. Additionally, due to the high level of research activities in this area, the exchange between experts is important in order to share and discuss new findings with the community at an early stage.

**Best Practices for Quantum Algorithm Development.** The development of new algorithms requires in-depth knowledge and expertise in theory and practice. Documented, reusable best practices for recurring problems, i.e. patterns, can support and guide people in the development of new quantum algorithms.

**Decision-Support for Quantum Applications and Vendors.** A two-stage decisionsupport is required to identify appropriate solutions for real-world use cases. First, quantum algorithms that prove to provide a solution for a given problem have to be identified. Second, the appropriate implementation and quantum hardware have to be selected for integration and execution. For the second stage the resource consumption of algorithms and implementations on different quantum hardware are of main interest (see section 4.2).

**Vendor-Agnostic Usage of Quantum Hardware.** Currently, various algorithm implementations from different vendors are available via proprietary SDKs that have been developed specifically for their hardware. To avoid vendor lock-in the quantum algorithm must be portable between different vendors which can be achieved by a standardized quantum programming language (see section 3.1 and 4.2).

**Data Transformation for Quantum Algorithms.** Especially for machine learning and artificial intelligence data of sufficient quality is essential. This applies to both, classical and quantum algorithms. Such data have to be made available and respectively encoded for the quantum algorithm [49].

**Quantum Application as a Service (QaaS).** The hybrid architecture of quantum applications consisting of classical and quantum parts increases the complexity of their deployment. Quantum applications provided "as a Service" via a self-service portal ease the utilization of the new technology (see section 3.2).

#### 6.3 Architecture

In Figure 19 the architecture of the collaborative quantum software platform is depicted. In essence, the platform consists of two parts: The analysis and development platform as depicted on the left of the figure for collecting, discussing, analyzing, and sharing knowledge, and the marketplace as depicted on the right that offers solutions in the form of quantum applications and consulting services.

The analysis and development platform addresses the needs of specialists and researchers in the field of quantum computing and software engineering. In a first step, knowledge in the form of publications, software artifacts, datasets, or web content can be placed on the platform – either manually via a user interface or automatically using a crawler. This knowledge can originate from various sources, such as arXiv.org or github.com. In a first step it can be stored as raw data in the QAlgo & data content store. Content of interest has to be extracted from these raw data, such as a quantum algorithm described in a journal article. To facilitate collaboration among different disciplines and to create a common understanding, the representation of quantum circuits and mathematical expressions must be normalized. A qualified description of the knowledge artifact with metadata is also essential to find and link relevant knowledge. Therefore, metadata formats must be normalized and enriched. The knowledge artifacts are then stored and provided via an expert portal to specialists and scientists and via a customer portal to users looking for solutions for their use cases and the community of interested people.

Specialists and scientists can discuss, evaluate, and improve the different key entities on the platform. Algorithms and their implementations can be linked and evaluated based on defined metrics using the NISQ Analyzer (see section 4.3). Identified best practices, e.g., for creating entanglement, can then be stored as quantum patterns in a Quantum Computing Pattern Repository. These patterns ease the development of new algorithms as they provide proven solutions for frequently occurring problems at the design of quantum algorithms. Patterns solving specific problems can then be combine and applied for realizing a broader use case [20-22]. However, best practices are not only relevant for the development, but also for data preparation as input for quantum algorithms and the integration of quantum algorithms with classical applications. Data preparation is essential, and must especially be considered in the NISQ era.

Since most quantum algorithms are hybrid algorithms, execution of quantum applications means a distributed deployment of hybrid quantum applications among classical and quantum hardware. Such applications can be stored for reuse in the Hybrid-App-Repository. For the quantum part, the quantum computer vendor and more specific a single QPU has to be selected, depending on the QPU properties, the algorithm implementation, and the input data. The platform automates this selection and provides a vendor-agnostic access to quantum hardware. For the deployment, technologies for classical computing are evaluated to provide an integrated deployment automation toolchain. Standards such as the Topology and Orchestration Specification for Cloud Applications (TOSCA) [56] have been developed precisely for this purpose to enable portability, interoperability, and the distribution across different environments [68, 69].



Thus, TOSCA as an international standard offers good foundation for an integration of classical and quantum deployment.

Fig. 19. Architecture for a collaborative quantum software platform ([44] Fig. 10)

While the expert portal is tailored to provide a sufficient user interface and toolchain addressing the needs of quantum computing experts the marketplace on the right of Figure 19 enables service providers and further stakeholders, such as consultants, to offer solutions. Customers can place requests for solutions for certain problems or use cases at hand. It is further intended to also allow consulting services to be offered in addition to hybrid quantum applications and their deployments. This means that also business models besides the development and distribution are enabled by the interplay of the marketplace and the analysis and development platform. For example, hybrid quantum applications can be provided as a service, which is enabled through the automated deployment capabilities by means of a TOSCA orchestrator such as Open-TOSCA [12, 57] or Cloudify [15]. Further, the selection of quantum algorithms fitting to specific constraints of quantum hardware can be supported by the NISQ Analyzer and the discussions of experts. With the help of the marketplace, knowledge and software artifacts, such as quantum algorithm implementations and hybrid quantum applications, can be monetized. Every turnover on the platform leads to incentives for participating experts to make further knowledge available on the platform.

#### 6.4 Sample Research Questions

The platform provides the basis for the technical realization of the research questions already discussed. However, further questions are raised: What are best practices for data preparation as input for quantum algorithms? What are best practices for integrating quantum algorithms with classical applications? How to combine the best practices

in quantum computing with other domains such as cloud computing? Which metadata is required to adequately describe the key entities on the platform?

## 7 Conclusions

New possibilities to solve classically intractable problems based on quantum computing is at the horizon. Quantum computers appear as part of the cloud infrastructure, and based on the hybrid nature of quantum-based applications, cloud computing techniques will contribute to the discipline of building them. Lots of new research questions appeared.

To evaluate promising application areas, we presented a use case from quantum humanities in which quantum machine learning techniques were used to derive insights from a comprehensive data set.

We are about to build the collaborative quantum application platform, and exploit it for several use cases, especially in the area of machine learning. A pattern language for quantum computing as well as tooling support [42, 84] are under construction. Research on the removal of hardware dependencies including deployment of hybrid quantum applications is ongoing.

Acknowledgements. We are grateful to Marie Salm and Manuela Weigold for discussing several subjects of this paper. Also, our thanks go to Philipp Wundrack, Marcel Messer, Daniel Fink and Tino Strehl for their valuable input and implementing several aspects of our use case.

This work was partially funded by the BMWi project PlanQK (01MK20005N) and the Terra Incognita project Quantum Humanities funded by the University of Stuttgart.

## References

- Aimeur, E., Brassard, G., Gambs, S.: Quantum Clustering Algorithms. In Proceedings of the 24th International Conference on Machine Learning, Corvallis, OR (2007)
- Amin, M. A., Andriyash, E., Rolfe, J., Kulchytskyy, B., Melko, R.: Quantum Boltzmann Machine. In Phys. Rev. X 8, 021050 (2018)
- Ankerst, M, Breunig, M. M. Kriegel, H, Sander, J.: OPTICS: Ordering Points To Identify the Clustering Structure. In: ACM SIGMOD international conference on Management of data. ACM Press (1999)
- Barzen, J.: Wenn Kostüme sprechen Musterforschung in den Digital Humanities am Beispiel vestimentärer Kommunikation im Film (in German). Dissertation University Cologne (2018)
- Barzen, J.: Taxonomien kostümrelevanter Parameter: Annäherung an eine Ontologisierung der Domäne des Filmkostüms (in German). Technical Report University Stuttgart No. 2013/04 (2013)
- Barzen, J., Breitenbücher, U., Eusterbrock, L., Falkenthal, M., Hentschel, F., Leymann, F.: The vision for MUSE4Music. Applying the MUSE method in musicology. In Comput Sci Res Dev, Proceedings of SummerSoC 2016, 32 (2016). https://doi.org/10.1007/s00450-016-0336-1

- Barzen, J., Falkenthal, M., Leymann, F.: Wenn Kostüme sprechen könnten: MUSE Ein musterbasierter Ansatz an die vestimentäre Kommunikation im Film (in German). In Bockwinkel, P., Nickel, B., Viehhauser, G., (ed.) Digital Humanities. Perspektiven der Praxis, Frank & Timme (2018)
- Barzen, J., Leymann, F.: Quantum Humanities: A First Use Case for Quantum-ML in Media Science. In ISAAI'19 Proceedings — Artificial Intelligence, Digitale Welt. Vol. 4(1) (2020).
- Barzen, J., Leymann, F.: Quantum humanities: a vision for quantum computing in digital humanities. In SICS Softw.-Inensiv. Cyber-Phys. Syst., Proceedings of SummerSoC 2019 (2019). https://doi.org/10.1007/s00450-019-00419-4
- 10. Berry, D., (ed.): Understanding Digital Humanities. Palgrave (2012)
- Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., Lloyd, S.: Quantum Machine Learning. In Nature 549, 195-202 (2017)
- Binz, T., Breitenbücher, U., Haupt, F., Kopp, O., Leymann, F., Nowak, A., Wagner, S.: OpenTOSCA – A Runtime for TOSCA-based Cloud Applications. In Proceedings of the 11th International Conference on Service-Oriented Computing (ICSOC 2013), pages 692-695. Springer (2013)
- 13. Booth Jr, J.: Quantum Compiler Optimizations (2012). arXiv:1206.3348
- Ciliberto, C., Herbster, M., Ialongo, A.D., Pontil M., Rocchetto, A., Severini, S., Wossnig, L.: Quantum machine learning: a classical perspective. In Proc.R.Soc.A 474: 20170551 (2018)
- 15. Cloudify (2020). https://cloudify.co/ (accessed 07/09/2020)
- 16. Cox, T.F., Cox, M.A.A.: Multidimensional Scaling. Chapman and Hall (2001)
- Crooks, G. E.: Performance of the Quantum Approximate Optimization Algorithm on the Maximum Cut Problem (2018). arXiv:1811.08419v1
- Dunjko, V., Taylor, J. M., Briegel, H. J.: Quantum-Enhanced Machine Learning. In Phys. Rev. Lett. 117, 130501 (2016)
- Falkenthal, M., Barzen, J., Breitenbücher, U., Brügmann, S., Joos, D., Leymann, F., Wurster, M.: Pattern Research in the Digital Humanities: How Data Mining Techniques Support the Identification of Costume Patterns. In Comput Sci Res Dev, Proceedings of SummerSoC 2016, 32 (2016). https://doi.org/10.1007/s00450-016-0336-1
- Falkenthal, M., Barzen, J., Breitenbücher, U., Fehling, C., Leymann, F.: Effective Pattern Application: Validating the Concept of Solution Implementation in Different Domains. In International Journal on Advances in Software, 7(3&4):710-726 (2014)
- Falkenthal, M., Barzen, J., Breitenbücher, U., Fehling, C., Leymann, F., Hadjakos, A., Hentschel, F., Schulze, H.: Leveraging Pattern Applications via Pattern Refinement. In Proceedings of the Inernational Conference on Pursuit of Pattern Languages for Social Change (PURPLSOC), pages 38-61, epubli GmbH (2016)
- Falkenthal, M., Barzen, J., Breitenbücher, U., Leymann, F.: Solution Languages: Easing Pattern Composition in Different Domains. In International Journal on Advances in Software, 10 (3&4):263-274 (2017)
- Falkenthal, M., Barzen, J., Dörner, S., Elkind, V., Fauser, J., Leymann, F., Strehl, T.: Datenanalyse in den Digital Humanities – Eine Annäherung an Kostümmuster mittels OLAP Cubes (in German). In Datenbanksysteme für Business, Technologie und Web (BTW), 16. Fachtagung des GI-Fachbereichs "Datenbanken und Informationssysteme" (2015)
- Farhi, E., Goldstone, J., Gutmann, S.: A Quantum Approximate Optimization Algorithm. In MIT-CTP/4610 (2014)
- Ghobadi, R., Oberoi, J. S., Zahedinejhad, E.: The Power of One Qubit in Machine Learning (2019). arXiv:1905.01390v2

- Grimsley, H. R., Economou, S. E., Barnes, E., Mayhall, N. J.: An adaptive variational algorithm for exact molecular simulations on a quantum computer. In Nat Commun 10, 3007 (2019)
- Guerreschi, G. G., Smelyanskiy, M.: Practical optimization for hybrid quantum-classical algorithms (2017). arXiv:1701.01450v1
- Havenstein, Ch., Thomas, D., Chandrasekaran, S.: Comparisons of Performance between Quantum and Classical Machine Learning. In SMU Data Science Review: Vol. 1: No. 4, (2018)
- Häner, T., Steiger, D. S., Svore, K., Troyer, M.: A software methodology for compiling quantum programs. Quantum Science and Technology, 3(2):020501 (2018)
- Heyfron, L. E., Campbell, E. T.: An efficient quantum compiler that reduces T count. Quantum Science and Technology, 4(1):015004 (2018)
- Hinton, G.E., Salakhutdinov, R. R.: Reducing the dimensionality of data with neural networks. In Science 313 (5786) (2006)
- Hinton, G.E.: A Practical Guide to Training Restricted Boltzmann Machines. In Montavon G., Orr G.B., Müller KR. (eds) Neural Networks: Tricks of the Trade. Lecture Notes in Computer Science, vol 7700. Springer (2012)
- Hofmann, Th., Schölkopf, B., Smola, A. J.: Kernel Methods in Machine Learning. In The Annals of Statistics 2008, Vol. 36, No. 3, 1 (2008)
- Itoko, T., Raymond, R., Imamichi, T., Matsuo, A.: Optimization of quantum circuit mapping using gate transformation and commutation. Integration, 70:43–50 (2020)
- Javadi-Abhari, A., Nation, P., Gambetta, J.: Qiskit write once, target multiple architectures (2019). https://www.ibm.com/blogs/research/2019/11/qiskit-for-multiple-architectures/ (accessed 07/09/2020)
- Javadi-Abhari, A., Patil, S., Kudrow, D., Heckey, J., Lvov, A., Chong, F. T., Martonosi, M.: ScaffCC: Scalable Compilation and Analysis of Quantum Programs. Parallel Computing, 45:2–17 (2015)
- Jozsa, R., Linden, N.: On the role of entanglement in quantum-computational speed-up. Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences, 459(2036):2011–2032 (2003)
- Laflamme, R., Miquel, C., Paz, J. P., Zurek, W. H.: Perfect Quantum Error Correcting Code. Physical Review Letters, 77(1):198 (1996)
- LaRose, M.: Overview and Comparison of Gate Level Quantum Software Platforms (2019). arXiv:1807.02500v2
- Leymann, F.: Towards a Pattern Language for Quantum Algorithms. In First International Workshop on Quantum Technology and Optimization Problems (QTOP 2019), Springer International Publishing (2019)
- Leymann, F., Barzen, J.: The bitter truth about gate-based quantum algorithms in the NISQ era. In: Quantum Sci. Technol. 5 044007 (2020)
- 42. Leymann, F., Barzen, J.; Pattern Atlas (2020). arXiv:2006.05120
- Leymann, F., Barzen, J., Falkenthal, M.: Towards a Platform for Sharing Quantum Software. In Proceedings of the 13<sup>th</sup> Advanced Summer School on Service Oriented Computing 2019. IBM Research Division (2019)
- Leymann, F., Barzen, J., Falkenthal, M., Vietz, D., Weder, B., Wild, K.: Quantum in the Cloud: Application Potentials and Research Opportunities. In: Proceedings of the 10<sup>th</sup> International Conference on Cloud Computing and Services Science (CLOSER 2020), pages 9-24, SciTePress (2020).
- Lloyd, S., Mohseni, M., Rebentrost, P.: Quantum principal component analysis. In Nature Physics 10, 631 (2014)

- Manning, Ch. D., Raghavan, P., Schütze, H.: An Introduction to Information Retrieval. Cambridge University Press (2009)
- Maslov, D., Dueck, G. W., Miller, D. M., Negrevergne, C.: Quantum Circuit Simplification and Level Compaction. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 27(3):436–444 (2008)
- McCaskey, A. J., Lyakh, D., Dumitrescu, E., Powers, S., Humble, T. S.: XACC: a systemlevel software infrastructure for heterogeneous quantum classical computing. Quantum Science and Technology (2020)
- 49. Mitarai, K., Kitagawa, M., Fujii, K.: Quantum analog-digital conversion. In Phys. Rev. A 99(1). American Physical Society (2019)
- Mohseni, M., Read, P., Neven, H.: Commercialize early quantum technologies. In Nature 543(7644) (2017)
- Montanaro, A.: Quantum algorithms: an overview. In npj Quantum Information 2, 15023 (2016)
- MUSE (2020). https://www.iaas.uni-stuttgart.de/forschung/projekte/muse/ (accessed 09/07/2020)
- MUSE GitHub (2020) https://github.com/Muster-Suchen-und-Erkennen/muse-docker (accessed 07/09/2020)
- National Academies of Sciences, Engineering, and Medicine: Quantum Computing: Progress and Prospects. The National Academies Press, Washington, DC (2019)
- Nielsen, M. A., Chuang, I. L.: Quantum Computation and Quantum Information. Cambridge University Press (2016)
- OASIS (2019) TOSCA Simple Profile in YAML Version 1.2. OASIS. (accessed 07/09/2020).
- 57. OpenTOSCA (2020). http://www.opentosca.org/ (accessed 07/09/2020)
- 58. Otterbach, J. S., Manenti, R., Alidoust, N., Bestwick, A., Block, M., Bloom, B., Caldwell, S., Didier, N., Schuyler Fried, E., Hong, S., Karalekas, P., Osborn, C. B., Papageorge, A., Peterson, E. C., Prawiroatmodjo, G., Rubin, N., Ryan, C. A., Scarabelli, D., Scheer, M., Sete, E. A., Sivarajah, P., Smith, R: S., Staley, A., Tezak, N., Zeng, W. J., Hudson, A., Johnson, B. R., Reagor, M., da Silva, M. P., Rigetti, C.: Unsupervised Machine Learning on a Hybrid Quantum Computer (2017). arXiv:1712.05771
- Peruzzo, A., McClean, J., Shadbolt, P., Yung, M.-H., Zhou, X.-Q., Love, P. J., Aspuru-Guzik, A., O'Brien, J.L.: A variational eigenvalue solver on a quantum processor. In Nature Communications, 5:4213, (2014)
- Plesch, M., Brukner, Č.: Quantum-state preparation with universal gate decompositions. In Phys. Rev. A 83, 032302, (2011)
- 61. Preskill, J.: Quantum Computing in the NISQ era and beyond. In Quantum 2, 79 (2018)
- 62. PyQuil (2020). https://github.com/rigetti/pyquil (accessed 07/09/2020)
- 63. PyTorch (2020). https://pytorch.org/ (accessed 07/09/2020)
- 64. Qiskit (2020). https://qiskit.org/ (accessed 07/09/2020)
- Rahaman, M., Islam, M. M.: A Review on Progress and Problems of Quantum Computing as aService (QcaaS) in the Perspective of Cloud Computing. Global Journal of Computer Science and Technology (2015)
- 66. Rebentrost, P., Mohseni, M., Lloyd, S.: Quantum support vector machine for big data classification. In Phys. Rev. Lett. 113, 130503 (2014)
- Romero, J., Olson, J. P., Aspuru-Guzik, A.: Quantum autoencoders for efficient compression of quantum data (2017). arXiv:1612.02806v2

- Saatkamp, K., Breitenbücher, B., Kopp, O., Leymann, F.: Method, formalization, and algorithms to split topology models for distributed cloud application deployments. Computing 102(2): 343-363 (2019)
- Saatkamp, K., Breitenbücher, U., Kopp, O., Leymann, F.: Topology Splitting and Matching für Multi-Cloud Deployments. In Proceedings of the 7<sup>th</sup> International Conference on Cloud Computing and Service Science (CLOSER 2017), pages 247-258, SciTePress (2017)
- Salm, M., Barzen, J., Breitenbücher, U., Leymann, F., Weder, B., Wild, K.: A Roadmap for Automating the Selection of Quantum Computers for Quantum Algorithms (2020). arXiv:2003.13409
- Scherer, A., Valiron, B., Mau, S.-C., Alexander, S., Van den Berg, E., and Chapuran, T. E.: Concrete resource analysis of the quantum linear-system algorithm used to compute the electromagnetic scattering cross section of a 2d target. Quantum Information Processing, 16(3):60 (2017)
- Schuld, M., Fingerhuth, M., Petruccione, F.: Implementing a distance-based classifier with a quantum interference circuit (2017). arXiv:1703.10793
- Schuld, M., Killoran, N.: Quantum machine learning in feature Hilbert spaces. In Phys. Rev. Lett. 122, 040504 (2019)
- Schuld, M., Sinayskiy, I., Petruccione, F.: Quantum Computing for Pattern Classification. In Trends in Artificial Intelligence, LNAI 8862, Springer (2014)
- Schuld, M., Sinayskiy, I., Petruccione, F.: An introduction to quantum machine learning. In Contemporary Physics, 56:2, 172-185 (2015)
- Shende, V. V., Markov, I. L.: Quantum Circuits For Incompletely Specified Two-Qubit Operators. In Quantum Information & Computation 5 (2005)
- Smith, R. S., Curtis, M. J., Zeng, W. J.: A practical quantum instruction set architecture (2016). arXiv:1608.03355.
- Steiger, D. S., Haner, T., and Troyer, M.: ProjectQ: An open source software framework for quantum computing. Quantum, 2:49 (2018)
- Suchara, M., Kubiatowicz, J., Faruque, A., Chong, F. T., Lai, C.-Y., and Paz, G.: QuRE: The quantum resource estimator toolbox. In 2013 IEEE 31st International Conference on Computer Design (ICCD), pages 419–426. IEEE (2013)
- Svore, K. M., Aho, A. V., Cross, A. W., Chuang, I., Markov, I. L.: A layered software architecture for quantum computing design tools. Computer, 39(1):74–83 (2006)
- Taherkhania, A., Cosmaa, G., McGinnity, T. M.: Deep-FS: A feature selection algorithm for Deep Boltzmann Machines. In Neurocomputing 322 (2018)
- 82. TensorFlow (2020). https://www.tensorflow.org/ (accessed 07/09/2020)
- 83. TensorFlow Quantum (2020). https://www.tensorflow.org/quantum (accessed 07/09/2020)
- Weigold, M., Barzen, J., Breitenbücher, U., Falkenthal, M., Leymann, F., Wild, K.: Pattern Views: Concept and Tooling for Interconnected Pattern Languages (2020). arXiv:2003.09127
- Wu, Z., Palmer, M.: Verb semantics and lexical selection. In Proceedings of the 32nd Annual Meeting of the Associations for Computational Linguistics, Las Cruces, New Mexico, (1994)
- Zhang, P., Li, S., Zhou, Y.: An Algorithm of Quantum Restricted Boltzmann Machine Network Based on Quantum Gates and Its Application. In Shock and Vibration Volume 2015, Article ID 756969 (2015). https://doi.org/10.1155/2015/756969
- Zhang, Y., Deng, H., Li, Q., Song, H., Nie, L.: Optimizing Quantum Programs against Decoherence: Delaying Qubits into Quantum Superposition. In 2019 International Symposium on Theoretical Aspects of Software Engineering (TASE), pages 184–191. IEEE (2019)

 Zhou, L., Wang, S.-T., Choi, S., Pichler, H., Lukin, M. D.: Quantum Approximate Optimization Algorithm: Performance, Mechanism, and Implementation on Near-Term Devices (2019). arXiv:1812.01041v2