

TensorCircuit: a Quantum Software Framework for the NISQ Era

Shi-Xin Zhang¹, Jonathan Allcock², Zhou-Quan Wan^{3,1}, Shuo Liu^{3,1}, Jiace Sun⁴, Hao Yu⁵, Xing-Han Yang^{6,1}, Jiezhong Qiu¹, Zhaofeng Ye¹, Yu-Qin Chen¹, Chee-Kong Lee⁷, Yi-Cong Zheng¹, Shao-Kai Jian⁸, Hong Yao³, Chang-Yu Hsieh¹, and Shengyu Zhang¹

¹Tencent Quantum Laboratory, Tencent, Shenzhen, Guangdong 518057, China

²Tencent Quantum Laboratory, Tencent, Hong Kong, China

³Institute for Advanced Study, Tsinghua University, Beijing 100084, China

⁴Division of Chemistry and Chemical Engineering, California Institute of Technology, Pasadena, CA 91125, USA

⁵Department of Electrical and Computer Engineering, McGill University, Quebec H3A 0E9, Canada

⁶Shenzhen Middle School, Shenzhen, Guangdong 518025, China

⁷Tencent America, Palo Alto, California 94306, USA

⁸Department of Physics, Brandeis University, Waltham, Massachusetts 02453, USA

TensorCircuit is an open source quantum circuit simulator based on tensor network contraction, designed for speed, flexibility and code efficiency. Written purely in Python, and built on top of industry-standard machine learning frameworks, TensorCircuit supports automatic differentiation, just-in-time compilation, vectorized parallelism and hardware acceleration. These features allow TensorCircuit to simulate larger and more complex quantum circuits than existing simulators, and are especially suited to variational algorithms based on parameterized quantum circuits. TensorCircuit enables orders of magnitude speedup for various quantum simulation tasks compared to other common quantum software, and can simulate up to 600 qubits with moderate circuit depth and low-dimensional connectivity. With its time and space efficiency, flexible and extensible architecture and compact, user-friendly API, TensorCircuit has been built to facilitate the design, simulation and analysis of quantum algorithms in the Noisy Intermediate-Scale Quantum (NISQ) era.

Contents

1	Introduction	2
1.1	Challenges in simulating quantum circuits	3
1.2	Machine learning libraries	4
1.3	The next phase of quantum software	4
2	TensorCircuit Overview	5
2.1	Design philosophy	6
2.2	Tensor network engine	6
2.3	Installing and contributing to TensorCircuit	7
3	Circuits and gates	7
3.1	Preliminaries	8
3.2	Basic circuits and outputs	8
3.3	Specifying the input state and composing circuits	11
3.4	Circuit transformation and visualization	11

Chang-Yu Hsieh: kimhsieh@tencent.com

Shengyu Zhang: shengyuzhang@tencent.com

4	Gradients, optimization and variational algorithms	11
4.1	Optimization via ML backends	12
4.2	Optimization via SciPy	13
5	Density matrices and mixed state evolution	13
5.1	Density matrix simulation with tc.DMCircuit	14
5.2	Monte Carlo simulation with tc.Circuit	15
5.2.1	Externalizing the randomness	15
6	Advanced features	16
6.1	Conditional measurements and post-selection	16
6.1.1	Conditional measurements	16
6.1.2	Post-selection	17
6.2	Pauli string expectation	17
6.2.1	Pauli structures and weights	18
6.2.2	Explicit loop with c.expectation_ps	18
6.2.3	Expectations of Hamiltonians via operator_expectation	18
6.2.4	vmap over Pauli structures	19
6.3	vmap and vectorized_value_and_grad	21
6.3.1	Batched input states	21
6.3.2	Batched circuits	22
6.3.3	Batched cost function evaluation	22
6.3.4	Batched Machine Learning	23
6.3.5	Batched VQE	24
6.3.6	Batched Monte Carlo trajectory noise simulation	25
6.4	QuOperator and QuVector	25
6.4.1	QuVector as the input state for the circuit	27
6.4.2	QuVector as the output state of the circuit	27
6.4.3	QuOperator as the operator to be measured	27
6.4.4	QuOperator as a quantum gate	27
6.5	Custom contraction settings	28
6.5.1	Customized contraction path finder	28
6.5.2	Subtree reconfiguration	29
6.6	Advanced automatic differentiation	30
7	Integrated examples	32
7.1	Molecular VQE with TensorCircuit and OpenFermion	32
7.2	Quantum machine learning	34
7.3	Demonstration of barren plateaus	35
7.4	Very large circuit simulation	36
8	Outlook and concluding remarks	38
	References	38

1 Introduction

The landscape of open-source and proprietary software for simulating quantum computers [1] has grown in recent years. While features and functionality vary between packages, across the board users now have many high-quality options available for constructing and simulating quantum circuits. However, to deepen our understanding of quantum algorithm performance, researchers increasingly need to simulate larger and more complex quantum circuits, and to optimize quantum circuits that may contain a large number of tunable parameters. In spite of the promising state of existing quantum software, there remain a number of challenges in running large-scale, complex simulations. Here we introduce TensorCircuit , a new open source tensor network based quantum circuit simulator built to address these challenges. Written in Python, and designed for speed, flexibility and ease-of-use, TensorCircuit is built on top of a

number of industry leading libraries. Via the TensorFlow [2], JAX [3] and PyTorch [4] machine learning libraries, convenient access is provided for automatic differentiation, just-in-time compilation, vectorized parallelism, and hardware acceleration. Fast tensor network contraction is enabled by the state-of-the-art cotengra [5, 6] package, which also gives users customizable control over the tensor network contraction process. These features enable efficient optimization of parameterized quantum circuits, allowing for more complex cases to be modelled. In addition, the TensorCircuit syntax aims to allow complicated tasks to be implemented with a minimal amount of code, saving time spent coding as well as in simulation.

1.1 Challenges in simulating quantum circuits

As fully fault-tolerant quantum computers capable of running large scale quantum algorithms may still be many years away, considerable research effort has been spent investigating the prospects for quantum advantage in the nearer term. Some algorithms for Noisy Intermediate-Scale Quantum (NISQ) [7] quantum computers aim to leverage classical computational power to supplement quantum computers, which may have only a limited number of error-prone qubits under control. In particular, hybrid quantum-classical algorithms [8, 9] such as the Variational Quantum Eigensolver (VQE) [10] and the Quantum Approximate Optimization Algorithm (QAOA) [11] are based around the concept of parameterized quantum circuits (PQC). These circuits, which contain quantum gates with tunable parameters (for instance, single-qubit gates with variable rotation angles), are embedded in a classical optimization loop. By optimizing the value of the parameters in the circuit, one aims to drive the output state of the quantum circuit towards the solution to a given problem.

In a prototypical VQE example, one wishes to find the ground state energy E_0 of a quantum system with Hamiltonian H . The output of a PQC is a quantum state $|\psi(\boldsymbol{\theta})\rangle$, where $\boldsymbol{\theta}$ is a vector of tunable parameters. This trial state – known as an ansatz – forms a guess for the ground state wavefunction. By performing appropriate measurements on $|\psi(\boldsymbol{\theta})\rangle$, one can estimate the expected energy $\langle H \rangle_{\boldsymbol{\theta}} = \langle \psi(\boldsymbol{\theta}) | H | \psi(\boldsymbol{\theta}) \rangle$. By minimizing $\langle H \rangle_{\boldsymbol{\theta}}$ with respect to the parameters $\boldsymbol{\theta}$, one obtains an upper bound estimate of the ground state energy:

$$E_0 \leq \min_{\boldsymbol{\theta}} \langle H \rangle_{\boldsymbol{\theta}}. \quad (1)$$

There are a number of issues that affect the efficacy and efficiency for this approach. Firstly, for an accurate estimate of E_0 , the ansatz $|\psi(\boldsymbol{\theta})\rangle$ should, for some values of $\boldsymbol{\theta}$, be a good approximation to the true ground state. Whether this is so depends on the nature of the problem and the complexity of the PQC from which the ansatz is constructed. On the one hand, simple ansätze – for instance, the “hardware efficient” ansatz of [12] – may be easier to implement on real or simulated quantum computers, but may either not be sufficiently accurate (for instance, due to lack of physically-relevant structure or short depth) or else suffer from other issues, such as so-called barren plateaus in parameter space [13] or local minima in the energy landscape [14]. On the other hand, more complex ansätze, such as the unitary coupled cluster (UCC) [15] approach proposed for quantum chemistry problems, may require quantum circuit complexities and depths beyond what can currently be implemented or simulated, and the associated circuits must then be truncated or simplified. The ability to simulate larger and deeper quantum circuits would enable a systematic investigation of larger classes of ansätze.

Secondly, for a given ansatz, the evaluation of $\langle H \rangle_{\boldsymbol{\theta}}$ can be an involved process. For instance, consider the case where H is an n -qubit Hamiltonian, which can be expressed as a weighted sum of tensor products of Pauli operators, i.e.,

$$H = \sum_{j=1}^K \alpha_j P_j, \quad (2)$$

where α_j are real coefficients, each Pauli string P_j is of the form $P_j = \sigma_{i_1} \otimes \sigma_{i_2} \otimes \dots \otimes \sigma_{i_n}$, and σ_{i_j} are single-qubit Pauli operators or the identity. In the most straightforward approach, estimating $\langle H \rangle_{\boldsymbol{\theta}}$ is performed by first estimating the expectation of each term $\langle P_j \rangle_{\boldsymbol{\theta}}$ and then adding together these individual contributions. If the Hamiltonian consists of many Pauli terms, ways of speeding up the evaluation of Eq. (2) – for instance, by exploiting efficient representations of H (e.g. as a sparse matrix or Matrix Product Operator (MPO) [16]), or being able to compute multiple terms in parallel – can have a large impact on the computation time.

Thirdly, the optimization problem in Eq. (1) is, in general, non-convex. Thus, finding a global minimum is in general computationally intractable. However, if gradients of potentially complicated expressions can be efficiently evaluated, one can use gradient descent methods to find a local minimum which may yield a decent approximate solution. By initializing the optimization from a number of different positions (i.e. different values of θ), multiple local minima can be obtained, improving the chances that one of these gives a good solution. If these multiple solutions can be optimized in parallel, one may achieve large efficiency gains.

1.2 Machine learning libraries

The challenges discussed in the previous section overlap to a large degree with problems faced in machine learning, and especially deep learning [17]. Fortunately, the need to tackle machine learning problems of increasing complexity and to deal with datasets of ever-larger size, has led to the development of impressive software, and many frameworks are now available that combine powerful features with easy-to-use syntax. In particular:

Fast gradients. At the heart of all advanced machine learning packages is the ability to perform automatic differentiation (AD) [18, 19], of which the backpropagation algorithm used to train neural networks is a special case. AD enables efficient computation of the gradients of functions defined in code, and is vital to the optimization of many machine learning models.

JIT. Just-in-time compilation is a way of compiling certain parts of code during program execution. For interpreted languages such as Python, "jitting" a function can lead to large performance gains, with the time required to execute a jitted function often only a small fraction of the time needed if the function were interpreted. While the first time the function is called there may be an overhead cost required for compilation (staging time), this cost can be negligible to the time saved if the function is subsequently called many times.

Vectorization (VMAP). This feature allows a function to be evaluated on multiple points in parallel, with significant speedup compared with using a naïve for loop. In machine learning this allows one, for instance, to perform computations on batches of data at the same time.

Hardware acceleration. For complex machine learning models, the ability to execute code on multiple CPUs, GPUs and TPUs may be necessary for training to complete in a reasonable amount of time.

These powerful features are also beneficial in simulating quantum computers, and are particularly suited to variational quantum algorithms. Here, fast gradient evaluation via automatic differentiation gives simulators an inherent advantage over real quantum computers, which must estimate gradients in a less direct way; for instance by sampling the outputs for various input parameter choices and computing finite differences [20, 21]. Vectorization can (among other things) be used to evaluate PQC circuits with multiple parameter choices concurrently, or compute expectations of multiple Pauli strings simultaneously, and JIT and hardware acceleration provide for further time savings.

In addition, there is increasing interest in solving machine learning problems via quantum computing, as well as combining classical and quantum machine learning (QML) algorithms. Both of these can be facilitated by better integration between classical ML frameworks and quantum circuit simulators, and great value can be derived from a seamless integration of the two.

1.3 The next phase of quantum software

While quantum software has progressed a great deal in the last few years – with packages such as Qiskit [22], Cirq [23], ProjectQ [24], HiQ [25], Q# [26], Qibo [27], and qulacs [28] all offering powerful functionality and features – there remain significant advantages to be gained from efficient quantum simulation software supplemented with the power and features of state-of-the-art machine learning. Recently, a new generation of quantum software has started to emerge, with TensorFlow Quantum [29], PennyLane [30], Paddle Quantum [31] and MindQuantum [32] in the Python ecosystem making inroads here, and providing these features to varying degrees. However, to date, none of these fully combine all

	AD	JIT	VMAP	GPU	TN
Qiskit/Cirq/ProjectQ/Qulacs	.	.	.	✓✓	.
Qibo	✓	✓✓	.	✓✓	.
TensorFlow Quantum	✓	.	✓✓	.	.
PennyLane	✓✓	✓	✓	✓✓	.
TensorCircuit	✓✓	✓✓	✓✓	✓✓	✓✓

Table 1: Quantum software support for main machine learning paradigms and tensor network engine. Check marks qualitatively indicate the level of support for the corresponding features: ✓✓ = good support, ✓ = limited support, . = not supported. TN is for tensor network simulation engine support. Many well-known quantum software frameworks such as Qiskit or Cirq have no support for AD, JIT and VMAP; Qibo does not support AD with its most optimized and efficient backend; TensorFlow Quantum supports AD of expectation values but not directly of the wavefunction, and does not support quantum circuit simulation on GPU; PennyLane does not support VMAP of trainable parameters and JIT support is fragile as some methods are not JIT-compatible, etc.

	AD	JIT	VMAP	GPU
TensorFlow	✓✓	✓✓	✓✓	✓✓
JAX	✓✓	✓✓	✓✓	✓✓
PyTorch	✓✓	✓	✓	✓✓
NumPy

Table 2: Backends supported by TensorCircuit. Check marks qualitatively indicate the level of support for the corresponding features: ✓✓ = good support, ✓ = limited support, . = not supported. TensorFlow and JAX both offer comprehensive support for AD, JIT, VMAP and hardware acceleration, and are recommended for most tasks. PyTorch currently has limited support for JIT and VMAP as these are implemented in experimental modules and are not yet stable. In TensorCircuit we supplement the AD infrastructure of TensorFlow with a VMAP-compatible implementation of Jacobian and Hessian calculations, and enrich the functionality of TensorFlow’s VMAP, which in TensorCircuit now supports VMAP over multiple arguments. These improvements can be accessed via the unified backend abstraction provided by TensorCircuit. If no backend is chosen, TensorCircuit defaults to using NumPy as a backend, which does not support these advanced ML features. For more on the choice of ML backend, please see the documentation: [faq.html#which-ml-framework-backend-should-i-use](#).

of the key features of the previous section with fast quantum circuit simulation. That is, there is no good solution for space and time efficient noiseless and noisy quantum simulation that combines a tensor network engine with the machine learning paradigms of AD, JIT and VMAP, as well as GPU support. (See Table 1 for comparison of quantum software in terms of machine learning paradigm compatibility.) TensorCircuit has been designed to fill this gap, and give users a faster, more flexible and more convenient way to simulate quantum circuits and quantum processes.

2 TensorCircuit Overview

Our goal with TensorCircuit is to provide the first efficient, tensor network based quantum simulator that is fully compatible with the key features of modern machine learning frameworks, especially the programming paradigms of automatic differentiation, vectorized parallelism and just-in-time compilation. These features are provided via a number of popular machine learning backends which, at the time of writing, are TensorFlow, JAX and PyTorch (see Table. 2).

Integration with these backends allows for general hybrid quantum-classical models, where the outputs of a parameterized quantum circuit may be fed into a classical neural network or vice versa, and simulated seamlessly (see Figure 1). This integration is also key for research related to quantum machine learning [33].

Currently, there are very limited options for tensor network based quantum simulators, with most popular software making use of state vector simulators. State vector simulators are strongly limited by memory as wavefunction amplitudes are stored in full, and can thus struggle to simulate circuits with larger numbers of qubits. On the other hand, quantum circuits with large numbers of (possibly noisy) qubits but relatively short circuit depths – such as those corresponding to NISQ devices, whose qubits have short coherence times – fall into the applicable region of tensor network simulators.

2.1 Design philosophy

With seamless integration with modern machine learning paradigms supported by an efficient tensor network based simulation engine, the design of TensorCircuit has, from the start, been based on the following principles.

Speed. TensorCircuit uses tensor network contraction to simulate quantum circuits, and is compatible with state-of-the-art third party contraction engines. In contrast, the majority of current popular quantum simulators are state-vector based. The tensor contraction framework allows TensorCircuit, in many cases, to simulate quantum circuits with improved efficiency in terms of time and space compared with other simulators. JIT, AD, VMAP and GPU support can all also provide significant acceleration (often, of several orders of magnitude) in many scenarios.

Flexibility. TensorCircuit is designed to be backend agnostic, making it easy to switch between any of the machine learning backends with no change in syntax or functionality. Via different ML backends, there is flexibility to simulate hybrid quantum-classical neural networks, run code on CPUs, GPUs and TPUs, and switch between 32 bit and 64 bit precision data types.

Code efficiency. Modern machine learning frameworks such as TensorFlow, PyTorch and JAX have user-friendly syntax, allowing for powerful tasks to be carried out with a minimal amount of code. With TensorCircuit, we are similarly focused on a compact and easy-to-use API to boost readability and productivity. Compared with other popular quantum simulation software, TensorCircuit can often perform similar tasks with significantly less code (see [tc vs. tfq for VQE](#) for an example). The backend agnostic syntax of TensorCircuit additionally makes it easy to switch between ML frameworks with a single line of setup code.

Community focused. TensorCircuit is open source software, and we care about readability, maintainability and extensibility of the codebase. We invite all members of the quantum computing community to take part in its continued development and use.

2.2 Tensor network engine

TensorCircuit simulates quantum circuits using the tensor network formalism, which has a long history in computational physics and was more recently pioneered for quantum circuit simulation [34]. Indeed, the graphical representations for quantum circuits and tensor networks are consistent, rendering a direct and simple translation from quantum circuit simulation to tensor network contraction. In this picture, quantum circuits are represented by a network of low-rank tensors corresponding to individual quantum gates, and the computation of amplitudes, expectation values or other scalar quantities is performed by contracting the edges of the network until only a single node remains. The order, or path, in which the edges are contracted is important, and can have a large impact on the time and space required to contract the network [6, 35, 36]. See [37] for a good introduction to tensor networks from a physics perspective.

In the most general case, as for other types of quantum simulators, the time required to simulate quantum circuits via tensor network contraction is exponential in the number of qubits. However, in special cases – including many of practical relevance – tensor network contraction can offer significant advantages since it avoids the memory bottleneck that plagues full state simulators and, to date, the largest scale quantum computing simulations such as the simulation of random circuits used in quantum supremacy experiments [38, 39] have all been performed via this approach [40, 41, 42, 43, 44, 45].

The tensor network data structure used in TensorCircuit is powered by the TensorNetwork [46] package. In addition, TensorCircuit can utilize state-of-the-art external Python packages such as cotengra for selecting efficient contraction paths, and the contraction is then performed via `einsum` and `matmul` by the machine learning backend selected by the user (see Figure 2).

Architecture. The tensor network engine underlying the simulation of quantum circuits is built on top of various machine learning frameworks with an abstraction layer in between that unifies different backends. At the application layer, TensorCircuit also includes various advanced quantum algorithms based on our

latest research [47, 48, 49, 50, 51, 52, 53]. The overall software architecture is shown in Figure 3.

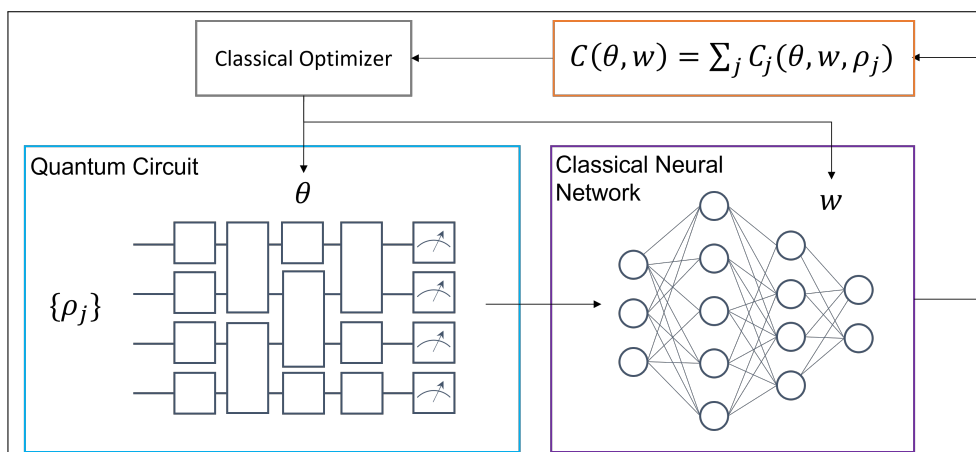


Figure 1: A general hybrid quantum-classical neural network, where a cost function C is summed over a batch of input states $\{\rho_j\}$ and is dependent on the parameters of a quantum circuit and classical neural network. By using a classical optimizer, the parameters of both networks can be iteratively improved. Integration with classical machine learning backends allows the entire end-to-end process to be seamlessly simulated in TensorCircuit, with vmap, jit and automatic differentiation enabling efficiency gains throughout the optimization process.




Figure 2: Schematic of the quantum circuit components in Figure 1. In TensorCircuit, the gates that comprise the quantum circuit are contained in a `tc.Circuit` object. Computation of circuit outputs is executed in two steps. First, a tensor contraction path is determined by a contraction engine, e.g., `cotengra`, and the backends take care of the actual contraction.

2.3 Installing and contributing to TensorCircuit

TensorCircuit is open-sourced under the Apache 2.0 license. The software is available on the Python Package Index (PyPI) and can be installed using the command `pip install tensorcircuit`.

The development of TensorCircuit is open-sourced and centered on GitHub: [TensorCircuit Repository](#). We welcome all members of the quantum community to contribute, whether it is

- Answering questions on the discussion page or issues page.
- Raising issues such as bug reports or feature requests on the issues page.
- Improving the documentation (docstrings/tutorials) by pull requests.
- Contributing to the codebase by pull requests.

For more details, please refer to the contribution guide: [contribution.html](#).

3 Circuits and gates

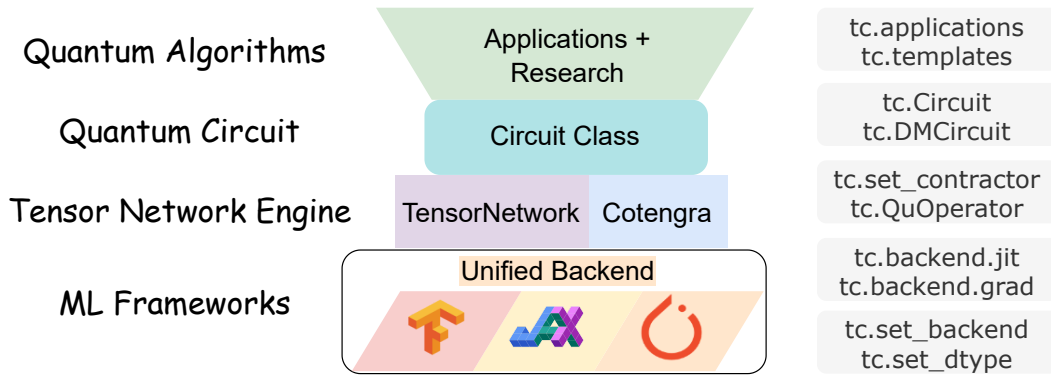


Figure 3: Software architecture of TensorCircuit . The abstraction layers are shown on the left while representative APIs are shown on the right. At the bottom, different machine learning frameworks such as TensorFlow, JAX and PyTorch are unified via a set of backend agnostic APIs. These backends, together with the tensor network engine, enable efficient quantum circuit simulation and implementation of quantum-classical hybrid algorithms and applications.

Jupyter notebook: [3-circuits-gates.ipynb](#)

In TensorCircuit , a quantum circuit on n qubits – which supports both noiseless and noisy simulations via Monte Carlo trajectory methods – is created by the `tc.Circuit(n)` API. Here we show how to create basic circuits, apply gates to them, and compute various outputs.

3.1 Preliminaries

In the remainder of this document, we assume that we have both TensorCircuit and NumPy imported as

```
1 import tensorcircuit as tc
2 import numpy as np
```

Furthermore we assume that a TensorCircuit backend has been set, e.g.

```
1 K = tc.set_backend("tensorflow")
```

and the symbol `K` that appears in code snippets (e.g. `K.real()`) refers to that backend. Other options for the `set_backend` method are “jax”, “pytorch” and “numpy” (the default backend).

In TensorCircuit , qubits are numbered from 0, with multiqubit registers numbered with the zeroth qubit on the left, e.g. $|0\rangle_{q_0} |1\rangle_{q_1} |0\rangle_{q_2}$. Unless needed, we will omit subscripts and use compact notation e.g. $|010\rangle$ to denote multiqubit states. X, Y, Z denote the standard single qubit Pauli operators, with subscripts e.g. X_3 to clarify which qubit is acted on. Expectation values of operators with respect to a state $|\psi\rangle$ such as $\langle\psi| Z \otimes I \otimes X |\psi\rangle$ will be denoted in shorthand as $\langle Z_0 X_2 \rangle$. Unless stated, expectation values are always with respect to the output state of a given quantum circuit. If that circuit is parameterized by a set of angles θ , then the parameter-dependent expectation value may be denoted by $\langle \cdot \rangle_{\theta}$.

In TensorCircuit the default runtime datatype is `complex64`, but if higher precision is required this can be set as follows:

```
1 tc.set_dtype("complex128")
```

3.2 Basic circuits and outputs

Consider the following two-qubit quantum circuit consisting of a Hadamard gate on qubit q_0 , a CNOT on qubits q_0 and q_1 , and a single qubit rotation $R_X(0.2)$ of qubit q_1 by angle 0.2 about the X axis (see Figure 4). Qubits are numbered from 0 with q_0 on the top row. This circuit can be implemented in TensorCircuit as

```
1 c = tc.Circuit(2)
2 c.h(0)
3 c.cnot(0, 1)
4 c.rx(1, theta=0.2)
```





Figure 4: A two-qubit circuit consisting of Hadamard, CNOT and single-qubit R_X rotation.

From this, various outputs can be computed.

Basic outputs. The full wavefunction can be obtained via

```
1 c.state()
```

which will output an array $[\alpha_{00}, \alpha_{01}, \alpha_{10}, \alpha_{11}]$ corresponding to the amplitudes of the $|00\rangle, |01\rangle, |10\rangle, |11\rangle$ basis states. The full wavefunction can also be used to generate the reduced density matrix of a subset of the qubits, e.g.

```
1 # reduced density matrix for qubit 1
2 s = c.state()
3 tc.quantum.reduced_density_matrix(s, cut=[0]) # cut: list of qubit indices to trace out
```

Amplitudes of individual basis vectors are computed by passing the corresponding bit-string value to the amplitude function. For example, the amplitude of the $|10\rangle$ basis vector is computed by

```
1 c.amplitude("10")
```

The unitary matrix corresponding to the entire quantum circuit can also be output:

```
1 c.matrix()
```

Measurements and samples. Random samples corresponding to Z -measurements, i.e. in the $\{|0\rangle, |1\rangle\}$ basis, on all qubits can be generated using

```
1 c.sample()
```

which will output a (bitstring, probability) tuple, comprising a binary string corresponding to the measurement outcomes of a Z measurement on all the qubits and the associated probability of obtaining that outcome. Z measurements on a subset of qubits can be performed with the `measure` command

```
1 # return (outcome, probability) of measuring qubit 0 in Z basis
2 print(c.measure(0, with_prob=True))
```

For measurement of multiple qubits, simply provide a list of indices to `measure`, e.g. if `c` were a 4-qubit circuit, measurement of qubits 1, 3 can be done via

```
1 c.measure(1, 3, with_prob=True)
```

Note that measurement gates do not need to be explicitly added to the circuit in order to compute these outcomes and the `measure` and `sample` commands do not collapse the circuit output state. Measurement gates can be added and used, for instance, when gates must be applied conditioned on mid-circuit measurement outcomes (see Section 6.1).

Expectation values. Expectation values such as $\langle X_0 \rangle$, $\langle X_1 + Z_1 \rangle$ or $\langle Z_0 Z_1 \rangle$ can be computed via the expectation method of a circuit object, where the operator is defined via `Gate` object or simply an array.

```
1 c.expectation([tc.gates.x(), [0]]) # <X0>
2 c.expectation([tc.gates.x() + tc.gates.z(), [1]]) # <X1 + Z1>
3 c.expectation([tc.gates.z(), [0]], [tc.gates.z(), [1]]) # <Z0 Z1>
```

and expectations of user-defined operators can also be computed by supplying the corresponding array of matrix elements. For instance, the operator $2X + 3Z$ can be expressed as a matrix as

$$\begin{pmatrix} 3 & 2 \\ 2 & -3 \end{pmatrix}$$

and implemented (assuming the observable is measured on qubit 0) as

```
1 c.expectation([np.array([[3, 2], [2, -3]]), [0]])
```

Expectations of Pauli strings. While expectations of products of Pauli operators, e.g. $\langle Z_0 X_1 \rangle$ can be computed using `c.expectation` as above, `TensorCircuit` provides another way of computing such expressions which may be more convenient for longer Pauli strings:

```
1 c.expectation_ps(x=[1], z=[0])
```

and longer Pauli strings can similarly be computed by providing lists of indices corresponding to the qubits that the X, Y, Z operators act on. For example, for an $n = 5$ qubit circuit, the expectation value $\langle Z_0 X_1 Z_2 Y_4 \rangle$ is computed as

```
1 c.expectation_ps(x=[1], y=[4], z=[0, 2])
```

Standard quantum gates. Beyond the CNOT, Hadamard and R_X gates we have encountered so far, `TensorCircuit` provides support for a wide variety of commonly encountered quantum gates. The full list of gates can be found by querying

```
1 tc.Circuit.sgates # non-parameterized gates
2 tc.Circuit.vgates # parameterized gates
```

The matrix corresponding to a given gate, e.g. the Hadamard `h` gate, can be accessed in the following way

```
1 tc.gates.matrix_for_gate(tc.gates.h())
```

Arbitrary unitaries. User-defined unitary gates may be implemented by specifying their matrix elements as an array. As an example, the unitary $S = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}$ – which can also directly be added by calling `c.s()` – can be implemented as

```
1 c.unitary(0, unitary = np.array([[1,0],[0,1j]]), name='S')
```

where the optional `name` argument specifies how this gate is displayed when the circuit is output to `LaTeX`.

Exponential gates. Gates of the form $e^{i\theta G}$ where matrix G satisfies $G^2 = I$ admit a fast implementation via the `exp1` command, e.g., the two qubit gate $e^{i\theta Z \otimes Z}$ acting on qubits 0 and 1

```
1 c.exp1(0, 1, theta=0.2, unitary=tc.gates._zz_matrix)
```

where `tc.gates._zz_matrix` creates a numpy array corresponding to the the matrix

$$Z \otimes Z = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

General exponential gates, where $G^2 \neq I$ can be implemented via the `exp` command:

```
1 c.exp(0, theta=0.2, unitary=np.array([[2, 0],[0, 1]]))
```

Non-unitary gates. `TensorCircuit` also supports the application of non-unitary gates by supplying a non-unitary matrix as the argument to `c.unitary`, e.g.

```
1 c.unitary(0, unitary=np.array([[1,2],[2,3]]), name='non_unitary')
```

Note that non-unitary gates will lead to an output state that is no longer normalized, since normalization is often unnecessary and takes additional computational time.

3.3 Specifying the input state and composing circuits

By default, quantum circuits are applied to the initial all-zero product state. Arbitrary initial states can be set by passing an array containing the input state amplitudes to the optional `inputs` argument of `tc.Circuit`. For example, the maximally entangled state $\frac{|00\rangle+|11\rangle}{\sqrt{2}}$ can be input as follows:

```
1 c1 = tc.Circuit(2, inputs=np.array([1, 0, 0, 1] / np.sqrt(2)))
```

Input states in Matrix Product State (MPS) form can also be input via the optional `mps_inputs` argument of `tc.Circuit`. See Section 6.4 for details.

Circuits that act on the same number of qubits can be composed together via the `c.append()` or `c.prepend()` commands. With `c1` defined as above, we can create a new circuit `c2` and then compose them together:

```
1 c2 = tc.Circuit(2)
2 c2.cnot(0, 1)
3
4 c3 = c1.append(c2)
```

This leads to a circuit C_3 which is equivalent to first applying C_1 and then C_2 .

3.4 Circuit transformation and visualization

`tc.Circuit` objects can be converted to and from Qiskit `QuantumCircuit` objects. Export to Qiskit is done by

```
1 c.to_qiskit()
```

and the resulting `QuantumCircuit` object can then be compiled and run on compatible physical quantum processors and simulators. Conversely, importing a `QuantumCircuit` object from Qiskit is done via

```
1 c = tc.Circuit.from_qiskit(QuantumCircuit)
```

There are two ways to visualize quantum circuits generated in `TensorCircuit`. The first is to use

```
1 print(c.tex())
```

which outputs the code for drawing the associated quantum circuit using the `LATEXquantikz` package [54]. The second method uses the `draw` function:

```
1 c.draw()
```

which is a shortcut for

```
1 qc = c.to_qiskit()
2 qc.draw()
```

Underlying circuit transformation and visualization utilities is the quantum intermediate representation (IR) for `TensorCircuit` objects, which can be obtained by

```
1 c.to_qir()
```

4 Gradients, optimization and variational algorithms

Jupyter notebook: [4-gradient-optimization.ipynb](#)

`TensorCircuit` is designed to make optimization of parameterized quantum gates easy, fast and convenient. Consider a variational circuit acting on n qubits, and consisting of k layers, where each layer comprises parameterized $e^{i\theta X \otimes X}$ gates between neighboring qubits followed by a sequence of single qubit parameterized Z and X rotations:

We now show how to implement such circuits in `TensorCircuit`, and how to use one of the machine learning backends to compute cost functions and gradients easily and efficiently. The circuit for general n, k and set of parameters can be defined as follows:




Figure 5: A parameterized, layered quantum circuit. Each gate is dependent on a separate parameter, here all schematically represented as θ .

```

1 def qcircuit(n, k, params):
2     c = tc.Circuit(n)
3     for j in range(k):
4         for i in range(n - 1):
5             c.expl(
6                 i, i + 1, theta=params[j * (3 * n - 1) + i], unitary=tc.gates._xx_matrix
7             )
8         for i in range(n):
9             c.rz(i, theta=params[j * (3 * n - 1) + n - 1 + i])
10            c.rx(i, theta=params[j * (3 * n - 1) + 2 * n - 1 + i])
11    return c

```

As an example, we take $n = 3, k = 2$, set TensorFlow as our backend, and define an energy cost function to be minimized

$$E = \langle X_0 X_1 \rangle_\theta + \langle X_1 X_2 \rangle_\theta.$$

```

1 n = 3
2 k = 2
3
4 K = tc.set_backend("tensorflow")
5
6
7 def energy(params):
8     c = qcircuit(n, k, params)
9     e = c.expectation_ps(x=[0, 1]) + c.expectation_ps(x=[1, 2])
10    return K.real(e)

```

K.grad and K.value_and_grad. Using the ML backend support for automatic differentiation, we can now quickly compute both the energy and the gradient of the energy (with respect to the parameters):

```

1 energy_val_grad = K.value_and_grad(energy)

```

This creates a function which, given a set of parameters as input, returns both the energy and the gradient of the energy. If only the gradient is desired, then this can be computed by `K.grad(energy)`. While we could run the above code directly on a set of parameters, if multiple evaluations of the energy will be performed, significant time savings can be had by using a just-in-time compiled version of the function:

```

1 energy_val_grad_jit = K.jit(energy_val_grad)

```

With `K.jit`, the initial evaluation of the energy and gradient may take longer, but subsequent evaluations will be noticeably faster than non-jitted code. We recommend always using `jit` as long as the function is ‘tensor-in, tensor-out’, and we have worked hard to make all aspects of the circuit simulator compatible with JIT.

4.1 Optimization via ML backends

With the energy function and gradients available, optimization of the parameters is straightforward. Below is an example of how to do this via stochastic gradient descent:

```

1 learning_rate = 2e-2
2 opt = K.optimizer(tf.keras.optimizers.SGD(learning_rate))
3
4
5 def grad_descent(params, i):
6     val, grad = energy_val_grad_jit(params)
7     params = opt.update(grad, params)
8     if i % 10 == 0:
9         print(f"i={i}, energy={val}")
10    return params
11
12
13 params = K.implicit_randn(k * (3 * n - 1))
14 for i in range(200):
15     params = grad_descent(params, i)

```

While this example was done with the TensorFlow backend, switching to JAX can be done easily. All that is required is to redefine the optimizer `opt` using the JAX optimization library `optax` [55]:

```

1 import optax
2 opt = tc.backend.optimizer(optax.sgd(learning_rate))

```

Then, choose JAX as the backend via

```

1 K = tc.set_backend("jax")

```

and perform the gradient descent exactly as above. Note that if no backend is set explicitly, TensorCircuit defaults to using NumPy as the backend, which does not allow for automatic differentiation.

4.2 Optimization via SciPy

An alternative to using the machine learning backends for the optimization is to use SciPy. This can be done via the `scipy_interface` API call, and allows for gradient based (e.g. BFGS) and non-gradient based (e.g. COBYLA) optimizers to be used, which are not available via the ML backends.

```

1 import scipy.optimize as optimize
2
3 f_scipy = tc.interfaces.scipy_interface(energy, shape=[k * (3 * n - 1)], jit=True)
4 params = K.implicit_randn(k * (3 * n - 1))
5 r = optimize.minimize(f_scipy, params, method="L-BFGS-B", jac=True)

```

The first line above specifies the shape of the parameters to be supplied to the function to be minimized, which here is the energy function. The `jit=True` argument automatically takes care of jitting the energy function. Gradient-free optimization can similarly be performed efficiently by supplying the `gradient=False` argument to `scipy_interface`:

```

1 f_scipy = tc.interfaces.scipy_interface(
2     energy, shape=[k * (3 * n - 1)], jit=True, gradient=False
3 )
4 params = K.implicit_randn(k * (3 * n - 1))
5 r = optimize.minimize(f_scipy, params, method="COBYLA")

```

5 Density matrices and mixed state evolution

Jupyter notebook: [5-density-matrix.ipynb](#)

TensorCircuit provides two methods for simulating noisy, mixed state quantum evolution. Full density matrix simulation of n qubits is provided by using `tc.DMCCircuit(n)`, and then adding quantum operations – both unitary gates as well as general quantum operations specified by Kraus operators – to the circuit. Relative to pure state simulation of n qubits via `tc.Circuit`, full density matrix simulation is twice as memory intensive, and thus the maximum system size simulatable will be half of what can be simulated in the pure state case. A less memory intensive option is to use the standard `tc.Circuit(n)` object and stochastically simulate open system evolution via Monte Carlo methods.

5.1 Density matrix simulation with tc.DMCCircuit

We illustrate this method below, by considering a simple circuit on a single qubit, which takes as input the mixed state corresponding to a probabilistic mixture of the $|0\rangle$ state and the maximally mixed state

$$\rho(\alpha) = \alpha|0\rangle\langle 0| + (1 - \alpha)I/2.$$

This state is then passed through a circuit which applies an X gate, followed by a quantum operation corresponding to an amplitude damping channel \mathcal{E}_γ with parameter γ . This has Kraus operators

$$K_0 = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-\gamma} \end{pmatrix}, \quad K_1 = \begin{pmatrix} 0 & \sqrt{\gamma} \\ 0 & 0 \end{pmatrix}$$

This circuit thus induces the evolution

$$\rho(\alpha) \xrightarrow{X} X\rho(\alpha)X \xrightarrow{\mathcal{E}_\gamma} \sum_{i=0}^1 K_i X\rho(\alpha)X K_i^\dagger$$

To simulate this in TensorCircuit, we first create a tc.DMCCircuit (density matrix circuit) object and set the input state using the `dminputs` optional argument (note that if a pure state input is provided to tc.DMCCircuit, this should be done via the `inputs` optional argument rather than `dminputs`). $\rho(\alpha)$ has the matrix form

$$\rho(\alpha) = \begin{pmatrix} \frac{1+\alpha}{2} & \\ & \frac{1-\alpha}{2} \end{pmatrix},$$

and thus (taking $\alpha = 0.6$) we initialize the density matrix circuit as follows

```
1 def rho(alpha):
2     return np.array([[ (1 + alpha) / 2, 0 ], [ 0, (1 - alpha) / 2 ]])
3
4
5 input_state = rho(0.6)
6 dmc = tc.DMCCircuit(1, dminputs=input_state)
```

Adding the X gate (and other unitary gates) is done in the same way as for pure state circuits:

```
1 dmc.x(0)
```

To implement a general quantum operation such as the amplitude damping channel, we use `general_kraus`, supplied with the corresponding list of Kraus operators.

```
1 def amp_damp_kraus(gamma):
2     K0 = np.array([[1, 0], [0, np.sqrt(1 - gamma)]])
3     K1 = np.array([[0, np.sqrt(gamma)], [0, 0]])
4     return K0, K1
5
6
7 K0, K1 = amp_damp_kraus(0.3)
8 dmc.general_kraus([K0, K1], 0) # apply channel with Kraus operators [K0,K1] to qubit 0
```

and the full density matrix output can be returned via

```
1 dmc.state()
```

In this example we input the Kraus operators for the amplitude damping channel manually, in order to illustrate the general approach to implementing quantum channels with the tc.DMCCircuit method. In fact, TensorCircuit includes built-in methods for returning the Kraus operators for a number of common channels, including the amplitude damping, depolarizing, phase damping and reset channels. For example, the Kraus operators for the phase damping channel with parameter γ

$$K_0 = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-\gamma} \end{pmatrix}, \quad K_1 = \begin{pmatrix} 0 & 0 \\ 0 & \sqrt{\gamma} \end{pmatrix}$$

can be returned by calling

```

1 gamma = 0.3
2 K0, K1 = tc.channels.phasedampingchannel(gamma)

```

and the phase damping channel added to a circuit via

```

1 dmc.general_kraus([K0, K1], 0)

```

The above operation can be further simplified using one API call:

```

1 dmc.phasedamping(0, gamma=0.3)

```

5.2 Monte Carlo simulation with tc.Circuit

Monte Carlo methods can be used to sample noisy quantum evolution using `tc.Circuit` instead of `tc.DMCCircuit`, where the mixed state is effectively simulated by an ensemble of pure states. As for density matrix simulation, quantum channels \mathcal{E} can be added to a circuit object by providing a list of their associated Kraus operators $\{K_i\}$. The API is the same as for the full density matrix simulation:

```

1 input_state = np.array([1, 1] / np.sqrt(2))
2 c = tc.Circuit(1, inputs=input_state)
3 c.general_kraus(tc.channels.phasedampingchannel(0.5), 0)
4 c.state()

```

In this framework though, the output of a channel acting on $|\psi\rangle$, i.e.

$$\mathcal{E}(|\psi\rangle\langle\psi|) = \sum_i K_i |\psi\rangle\langle\psi| K_i^\dagger$$

is viewed as an ensemble of states $\frac{K_i |\psi\rangle}{\sqrt{\langle\psi| K_i^\dagger K_i |\psi\rangle}}$ that each occurs with probability $p_i = \langle\psi| K_i^\dagger K_i |\psi\rangle$.

Thus, the code above stochastically produces the output of a single qubit initialized in state $|\psi\rangle = \frac{|0\rangle+|1\rangle}{\sqrt{2}}$ being passed through a phase damping channel with parameter $\gamma = 0.5$.

The Monte Carlo simulation of channels where the Kraus operators are all unitary matrices (up to a constant factor) can be more efficiently handled by using `unitary_kraus` instead of `general_kraus`. For instance, the depolarizing channel with Kraus operators parameterized by p_x, p_y, p_z :

$$K_0 = (1 - p_x - p_y - p_z) \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad K_1 = p_x \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad K_2 = p_y \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad K_3 = p_z \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

can be implemented by

```

1 px, py, pz = 0.1, 0.2, 0.3
2 c.unitary_kraus(tc.channels.depolarizingchannel(px, py, pz), 0)

```

where, in the second line, `tc.channel.depolarizingchannel(px, py, pz)` returns the required Kraus operators.

5.2.1 Externalizing the randomness

The `general_kraus` and `unitary_kraus` examples above both handle randomness generation from inside the respective methods. That is, when the list $[K_0, K_1, \dots, K_{m-1}]$ of Kraus operators is supplied to `general_kraus` or `unitary_kraus`, the method partitions the interval $[0, 1]$ into m contiguous intervals $[0, 1] = I_0 \cup I_1 \cup \dots \cup I_{m-1}$ where the length of I_i is proportional to the relative probability of obtaining the outcome i . Then a uniformly random variable x in $[0, 1]$ is drawn from within the method, and outcome i is selected based on which interval x lies in. In `TensorCircuit`, we have full backend agnostic infrastructure for random number generation and management. However, the interplay between `jit`, random numbers and backend switching is often subtle if we rely on the random number generation inside these methods. See [advance.html#randoms-jit-backend-agnostic-and-their-interplay](#) for details.

In some situations, it may be preferable to first generate the random variable from outside the method, and then pass the value generated into `general_kraus` or `unitary_kraus`. This can be done via the optional `status` argument:

```

1 px, py, pz = 0.1, 0.2, 0.3
2 x = K.implicit_randn()
3 c.unitary_kraus(tc.channels.depolarizingchannel(px, py, pz), 0, status=x)

```

This is useful, for instance, when one wishes to use `vmap` to batch compute multiple runs of a Monte Carlo simulation. This is illustrated in the example below, where `vmap` is used to compute 10 runs of the simulation in parallel:

```

1 def f(x):
2     c = tc.Circuit(1)
3     c.h(0)
4     c.unitary_kraus(tc.channels.depolarizingchannel(0.1, 0.2, 0.3), 0, status=x)
5     return c.state()
6
7
8 f_vmap = K.vmap(f, vectorized_argnums=0)
9 X = K.implicit_randn(10)
10 f_vmap(X)

```

Conceptually, the line

```
1 f_vmap = K.vmap(f, vectorized_argnums=0)
```

creates a function which acts as

$$f_{vmap} \begin{pmatrix} x_0 \\ x_1 \\ \vdots \end{pmatrix} = \begin{pmatrix} f(x_0) \\ f(x_1) \\ \vdots \end{pmatrix}$$

and the argument `vectorized_argnums=0` indicates that is the zeroth argument (in this case the only argument) of f that we wish to batch compute in parallel.

6 Advanced features

This section consists of advanced features that the general reader may skip on first reading. The reader interested in the benchmark studies of Section 7 may wish to refer to this section though, as those examples use a number of concepts from here.

6.1 Conditional measurements and post-selection

Jupyter notebook: [6-1-conditional-measurements-post-selection.ipynb](#)

TensorCircuit allows for two kinds of operations to be performed that are related to measurement outcomes. These are (i) conditional measurements, the outcomes of which can be used to control downstream conditional quantum gates, and (ii) post-selection, which allows the user to select the post-measurement state corresponding to a particular measurement outcome.

6.1.1 Conditional measurements

The `cond_measure` command is used to simulate the process of performing a Z measurement on a qubit, generating a measurement outcome with probability given by the Born rule, and collapsing the wavefunction in accordance with the measured outcome. The classical measurement outcome obtained can then act as a control for a subsequent quantum operation via the `conditional_gate` API and can be used, for instance, to implement the canonical teleportation circuit.

```

1 # quantum teleportation of state |psi> = a|0> + sqrt(1-a^2)|1>
2 a = 0.3
3 input_state = np.kron(np.array([a, np.sqrt(1 - a ** 2)]), np.array([1, 0, 0, 0]))
4
5 c = tc.Circuit(3, inputs=input_state)
6 c.h(2)
7 c.cnot(2, 1)
8 c.cnot(0, 1)
9 c.h(0)
10
11 # mid-circuit measurements

```



```

12 z = c.cond_measure(0)
13 x = c.cond_measure(1)
14
15 # if x = 0 apply I, if x = 1 apply X (to qubit 2)
16 c.conditional_gate(x, [tc.gates.i(), tc.gates.x()], 2)
17
18 # if z = 0 apply I, if z = 1 apply Z (to qubit 2)
19 c.conditional_gate(z, [tc.gates.i(), tc.gates.z()], 2)

```




Figure 6: Teleportation circuit implemented with `c.cond_measure` and `c.conditional_gate`.

6.1.2 Post-selection

Post-selection is enabled in `TensorCircuit` via the `post_select` method. This allows the user to select the post- Z -measurement state of a qubit via the `keep` argument. Unlike `cond_measure`, the state returned by `post_select` is not normalized. As an example, consider

```

1 c = tc.Circuit(2, inputs=np.array([1, 0, 0, 1] / np.sqrt(2)))
2 c.post_select(0, keep=1) # measure qubit 0, post-select on outcome 1
3 c.state()

```

which initializes a 2-qubit maximally entangled state $|\psi\rangle = \frac{|00\rangle + |11\rangle}{\sqrt{2}}$. The first qubit (q_0) is then measured in the Z -basis, and the unnormalized state $|11\rangle/\sqrt{2}$ corresponding to the measurement outcome 1 being post-selected.

This post-selection scheme with unnormalized states is fast and can, for instance, be used to explore various quantum algorithms and nontrivial quantum physics such as measurement-induced entanglement phase transitions [56, 57, 58, 59, 60].

6.2 Pauli string expectation

Jupyter notebook: [6-2-pauli-string-expectation.ipynb](#)

Minimizing the expectation values of sums of Pauli strings is a common task in quantum algorithms. For instance, in the VQE ground state preparation of an n -site transverse-field Ising model (TFIM) with Hamiltonian

$$H = \sum_{i=0}^{n-2} J_i X_i X_{i+1} - \sum_{i=0}^{n-1} h_i Z_i,$$

where the J_i, h_i are model parameters, one wishes to minimize

$$E(\theta) = \langle H \rangle_\theta := \sum_{i=0}^{n-2} J_i \langle X_i X_{i+1} \rangle_\theta - \sum_{i=0}^{n-1} h_i \langle Z_i \rangle_\theta \quad (3)$$

with respect to the circuit parameters θ . `TensorCircuit` provides a number of ways to compute expressions of this form, useful in different scenarios. At a high level these are

- Looping over terms, each computed by `c.expectation_ps`.
- Supplying a dense matrix, sparse matrix or MPO representation of the Hamiltonian to the `operator_expectation` function.

- Using `vmap` to compute each of the terms in a vectorized parallel manner, where each term is input as a *structure* vector.

Underlying these methods are a variety of ways of representing strings of Pauli operators and converting between these representations. Before going through the above methods in more detail, let us introduce the Pauli structure vector representation utilized in `TensorCircuit`.

6.2.1 Pauli structures and weights

A string of Pauli operators acting on n qubits can be represented as a length- n vector $v \in \{0, 1, 2, 3\}^n$, where the value of $v_i = j$ corresponds to σ_i^j , i.e. Pauli operator σ^j acting on qubit i (with $\sigma^0 = I, \sigma^1 = X, \sigma^2 = Y, \sigma^3 = Z$). For example, in this notation, if $n = 3$ the term $X_1 X_2$ corresponds to $v = [0, 1, 1]$. We refer to such a vector representation of a Pauli string as a *structure*, and a list of structures, one for each Pauli string term in the Hamiltonian, is used as the input to compute sums of expectation values in a number of ways.

```

1 # Pauli structures for Transverse Field Ising Model
2 structures = []
3 for i in range(n - 1):
4     s = [0 for _ in range(n)]
5     s[i] = 1
6     s[i + 1] = 1
7     structures.append(s)
8 for i in range(n):
9     s = [0 for _ in range(n)]
10    s[i] = 3
11    structures.append(s)

```

If each structure has an associated weight, e.g. the term $X_i X_{i+1}$ has weight J_i in Hamiltonian (3), then we define a corresponding tensor of weights

```

1 # Weights, taking J_i = 1.0, all h_i = -1.0
2 J_vec = [1.0 for _ in range(n - 1)]
3 h_vec = [-1.0 for _ in range(n)]
4 weights = tc.array_to_tensor(np.array(J_vec + h_vec))

```

6.2.2 Explicit loop with `c.expectation_ps`

As introduced in Section 3.2, given a `TensorCircuit` quantum circuit c , a single Pauli string expectation can be computed by supplying a list of indices to `c.expectation_ps`. The sum can then be computed by using a straightforward loop:

```

1 def tfim_energy(c, J_vec, h_vec):
2     e = 0.0
3     n = c._nqubits
4     for i in range(n):
5         e += h_vec[i] * c.expectation_ps(z=[i])
6     for i in range(n-1):
7         e += J_vec[i] * c.expectation_ps(x=[i, i+1])
8     return K.real(e)

```

6.2.3 Expectations of Hamiltonians via `operator_expectation`

Given a `TensorCircuit` quantum circuit c and an operator `op` representing the Hamiltonian, the expectation value of the energy can also be computed via the `operator_expectation` API from the `TensorCircuit` templates library

```

1 e = tc.templates.measurements.operator_expectation(c, op)

```

The operator `op` itself can be expressed in one of three forms: (i) dense matrix, (ii) sparse matrix, and (iii) Matrix Product Operator (MPO).

Dense Matrix Input. As a simple example, take the Hamiltonian

$$X_0 X_1 - Z_0 - Z_1.$$

This has dense matrix representation

$$\begin{pmatrix} -2 & & 1 \\ & 1 & \\ 1 & & 2 \end{pmatrix}$$

and the expectation value of the Hamiltonian (with respect to circuit `c`) can be computed by

```
1 op = tc.array_to_tensor([[ -2, 0, 0, 1], [0, 0, 1, 0], [0, 1, 0, 0], [1, 0, 0, 0]])
2 e = tc.templates.measurements.operator_expectation(c, op)
```

The matrix elements above were input by hand. TensorCircuit also provides a way of generating the matrix elements from the associated Pauli structures and weights, e.g.

```
1 structure = [[1, 1], [0, 3], [3, 0]]
2 weights = [1.0, -1.0, -1.0]
3 H_dense = tc.quantum.PauliStringSum2Dense(structure, weights)
```

Sparse Matrix Input. Significant computational advantage in terms of space and time can be obtained if the Hamiltonian is sparse, in which case a sparse representation of the operator is preferable. This can be implemented in a backend agnostic way by converting from a list of Pauli structures in a two-stage process. First we convert to a sparse numpy matrix in COO (COOrdinate) format. For example, with the structures and weights defined in Section 6.2.1, we call

```
1 H_sparse_numpy = tc.quantum.PauliStringSum2COO_numpy(structures, weights)
```

Then we can convert to a sparse tensor compatible with the selected backend `K`

```
1 H_sparse = K.coo_sparse_matrix(
2     np.transpose(np.stack([H_sparse_numpy.row, H_sparse_numpy.col])),
3     H_sparse_numpy.data,
4     shape=(2 ** n, 2 ** n),
5 )
```

and then call `operator_expectation` on this sparse tensor

```
1 e = tc.templates.measurements.operator_expectation(c, H_sparse)
```

MPO Input. The TFIM Hamiltonian, as a short-ranged spin Hamiltonian, admits an efficient Matrix Product Operator representation. Again this is a two-stage process using TensorCircuit. We first convert the Hamiltonian into an MPO representation via the TensorNetwork [46] or Quimb package [61]:

```
1 # generate the corresponding MPO by converting the MPO in tensornetwork package
2
3 Jx = np.array([1.0 for _ in range(n - 1)]) # strength of xx interaction (OBC)
4 Bz = np.array([1.0 for _ in range(n)]) # strength of transverse field
5 # Note the convention for the sign of Bz
6 hamiltonian_mpo = tn.matrixproductstates.mpo.FiniteTFI(Jx, Bz, dtype=np.complex64)
```

and then convert the MPO into a `QuOperator` object compatible with TensorCircuit.

```
1 hamiltonian_mpo = tc.quantum.tn2qop(hamiltonian_mpo) # QuOperator in TensorCircuit
```

The expectation value of the energy can then be computed via `operator_expectation`

```
1 e = tc.templates.measurements.operator_expectation(c, hamiltonian_mpo)
```

6.2.4 vmap over Pauli structures

Given a state $|s\rangle$, the expectation value $\langle s|P|s\rangle$ of a Pauli string P with a given structure can be computed as a function of tensor inputs as

```
1 # assume the following are defined
2 # state: 2**n vector of coefficients corresponding to input state
3 # structure: Pauli structure (i.e. list of integers {0,1,2,3}**n)
4
5 structure = tc.array_to_tensor(structure)
```

time (s)	CPU	GPU
explicit loop	65.7	119
vmap over Pauli structures	0.68	0.0018
dense matrix representation	0.26	0.0015
sparse matrix representation	0.008	0.0014

Table 3: Performance benchmarks for Pauli string sum evaluation methods, with Hamiltonian corresponding to an H₂O molecule. 12 qubits are used for the binary encoding of STO-3G orbitals. The qubit Hamiltonian contains 1390 Pauli string terms in total. Data was obtained using the JAX backend. GPU simulations were performed on the Nvidia T4 GPU, while CPU simulations used Intel(R) Core(TM) i7-9750H CPU @ 2.60GHz. The MPO representation is not applicable in this case, since the required bond dimension is too large.

```

6 state = tc.array_to_tensor(state)
7
8
9 def e(state, structure):
10     c = tc.Circuit(n, inputs=state)
11     return tc.templates.measurements.parameterized_measurements(
12         c, structure, onehot=True
13     )

```

where the `parameterized_measurements` function is used to compute the expectation of the output of a circuit. Then, if the Hamiltonian is represented by a list of Pauli structures $[v_1, \dots, v_k]$, `vmap` can be used to compute the expectation with respect to the circuit `c` of each term in parallel:

```

1 e_vmap = K.vmap(e, vectorized_argnums=1)

```

Similar to the example in Section 5.2.1, `vmapping` creates a function which acts as

$$e_{vmap} \left(s, \begin{pmatrix} \leftarrow v_1 \rightarrow \\ \vdots \\ \leftarrow v_k \rightarrow \end{pmatrix} \right) = \begin{pmatrix} e(s, v_1) \\ \vdots \\ e(s, v_k) \end{pmatrix}$$

i.e., outputs a vector of expectation values corresponding to terms in the Hamiltonian, and

```

1 vectorized_argnums=1

```

indicates that it is the first argument `v` of the `e(s,v)` function which should be computed in parallel, while the zeroth argument (i.e. `s`) is fixed. With structures and weights as defined in Section 6.2.1, the expectation value of the Hamiltonian with respect to the circuit `c` can then be computed as

```

1 s = c.state()
2 e_terms = e_vmap(s, structures)
3 hamiltonian_expectation = K.sum(e_terms * K.real(weights))

```

We benchmark these different approaches to Pauli string sum evaluation for Hamiltonians corresponding to an H₂O molecule and the TFIM spin model, with results in Tables 3 and 4, respectively. See [examples/vqeh2o_benchmark.py](#) and [examples/vqetfim_benchmark.py](#) for more details. In the H₂O case, we observe an acceleration of *85,000 times* for VQE evaluation using the sparse matrix representation compared to a naïve loop (also implemented in `TensorCircuit`) as used in most other quantum software.

General remark on benchmark times: Throughout this work, when averages times per circuit or gradient evaluation are reported (either using `TensorCircuit` or other software), the initial compilation time associated with JIT is not reported. In practice, we find that this initial compilation time is typically an order of ten to several tens larger than the time required for subsequent evaluations, and its effect on the total algorithmic running time depends on how many iterations it is amortized over. If a circuit is only evaluated a few times, the time required for JIT compilation may outweigh the benefits it brings. However, in typical variational quantum algorithms or in noise simulation based on Monte Carlo trajectories, circuits are evaluated a large number of times, and the amortized cost of JIT compilation is negligible.

time (s)	CPU	GPU
explicit loop	1.73	0.11
vmap over Pauli structures	10.68	0.20
sparse matrix representation	0.61	0.0086
MPO matrix representation	0.0007	0.0039

Table 4: Performance benchmark for different Pauli string sum evaluation on a 20-qubit TFIM Hamiltonian with open boundary conditions. The qubit Hamiltonian contains 39 Pauli string terms. Data was obtained using the JAX backend. GPU simulations used the Nvidia T4 GPU, while CPU simulations used Intel(R) Xeon(R) Platinum 8255C CPU @ 2.50GHz. The dense matrix representation is not applicable for systems of 20 qubits due to excessive memory requirements.

6.3 vmap and vectorized_value_and_grad

Jupyter notebook: [6-3-vmap.ipynb](#)

As we have seen in Sections 5.2.1 and 6.2.4, `vmap` allows for batches of function evaluations to be performed simultaneously in parallel. If batch evaluation of gradients as well as function values is required, then this can be done via `vectorized_value_and_grad`. In the simplest case, consider a function $f(x, y)$ where $x \in \mathbb{R}^p, y \in \mathbb{R}^q$ are both vectors, and one wishes to evaluate both $f(x, y)$ and $\sum_x \nabla_y f(x, y) = \sum_x \left(\frac{\partial f(x, y_1)}{\partial y_1}, \dots, \frac{\partial f(x, y_q)}{\partial y_q} \right)^\top$ over a batch x_1, x_2, \dots, x_k of inputs x . This is achieved by creating a new, vectorized value-and-gradient function

$$f_{vvg} \left(\begin{pmatrix} \leftarrow x_1 \rightarrow \\ \vdots \\ \leftarrow x_k \rightarrow \end{pmatrix}, y \right) = \left(\begin{pmatrix} f(x_1, y) \\ \vdots \\ f(x_k, y) \end{pmatrix}, \sum_{i=1}^k \nabla_y f(x_i, y) \right)$$

which takes as zeroth argument the batched inputs expressed as a $k \times p$ tensor, and as first argument the variables we wish to differentiate with respect to. The outputs are a vector of function values evaluated at all points (x_i, y) , and the gradient averaged over all those points. A toy example is implemented as follows:

```

1 def f(x, y):
2     return x[0] * x[1] * y[0] ** 2
3
4
5 f_vvg = K.vectorized_value_and_grad(f, argnums=1, vectorized_argnums=0)
6 X = tc.array_to_tensor([[1, 2], [2, 3], [0, -1]])
7 y = tc.array_to_tensor([2])
8 f_vvg(x_tensor, y)

```

`argnums` indicates which argument we wish to take derivatives with respect to, and `vectorized_argnums` indicates which argument corresponds to the batched input. It is even possible to set the values of `argnums` and `vectorized_argnums` to be the same, i.e., we batch compute over different initial values of the parameters we wish to optimize over. This can be useful, for example, in batched VQE computations (see Section 6.3.5).

6.3.1 Batched input states

Consider a quantum circuit $U(w)$ on n qubits, parameterized by weights $w = [w_1, \dots, w_k]$, and with state-dependent loss function $f(\psi, w)$, e.g.

$$\mathcal{L} = \sum_{\psi} f(\psi, w) = \sum_{\psi} \langle \psi | U^\dagger(w) Z_2 U(w) | \psi \rangle$$

Then, both the value $f(\psi, w)$ and the gradient with respect to the weights $\nabla_w f(\psi, w)$ can be evaluated for a batch of k input states $|\psi_1\rangle, \dots, |\psi_k\rangle$ simultaneously.

```

1 f_vvg = K.vectorized_value_and_grad(f, argnums=1, vectorized_argnums=0)
2 f_vvg(psi_matrix, weights)

```

The vectorized value and grad workflow for batched input states can be visualized in Figure 7.

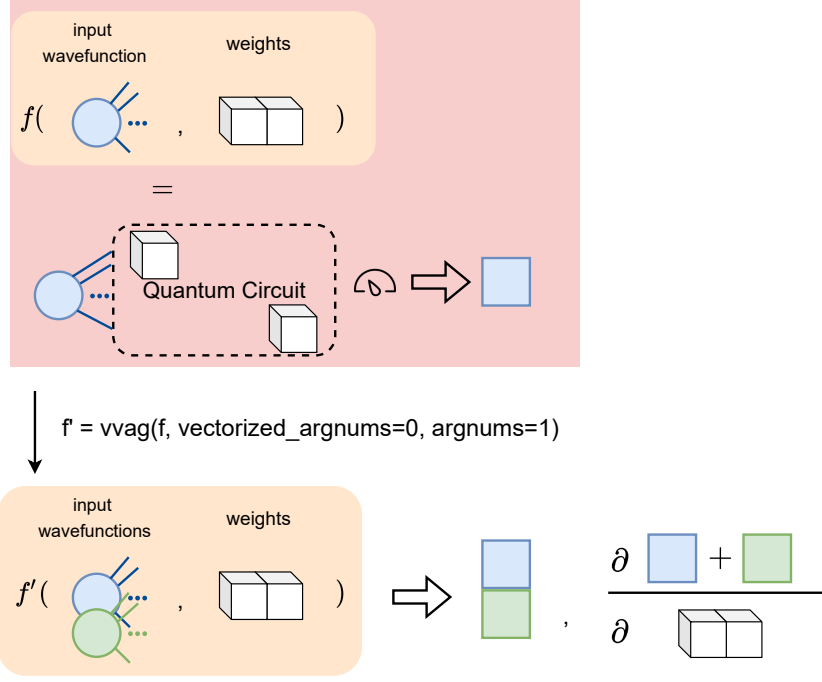


Figure 7: Schematic of applying `vectorized_value_and_grad` on quantum simulation tasks with multiple input states. In the figure, the function f' at the bottom is transformed from the original function f defined at the top via $f' = \text{K.vectorized_value_and_grad}(f)$. The wavefunction input is vectorized while differentiation is with respect to the weights input.

6.3.2 Batched circuits

Consider a family of parameterized quantum circuits $U_1(w), \dots, U_k(w)$ acting on the same number of qubits, where each circuit $U_i(w)$ can be expressed as a different parametrization of a single parent circuit, i.e.

$$U_j(w) = U(w, x_j)$$

where x_j is a vector of parameters. A loss function f is defined on these circuits, e.g. $f(w, x_j) = \langle 0 | U_j^\dagger X_1 U_j | 0 \rangle$ and its gradient with respect to the weights w can also be batch evaluated across all circuits simultaneously, e.g. by defining

$$f_{\text{vvg}} \left(w, \begin{pmatrix} \leftarrow x_1 \rightarrow \\ \vdots \\ \leftarrow x_k \rightarrow \end{pmatrix} \right) = \left(\begin{pmatrix} f(w, x_1) \\ \vdots \\ f(w, x_k) \end{pmatrix}, \sum_{i=1}^k \nabla_w f(w, x_i) \right)$$

In this case, the batched inputs correspond to the first argument, and the variables to differentiate with respect to correspond to the zeroth argument:

```

1 f_vvg = K.vectorized_value_and_grad(f, argnums=0, vectorized_argnums=1)
2 f_vvg(weights, x_matrix)

```

6.3.3 Batched cost function evaluation

As discussed in Section 6.2.4, `vmap` can be used to accelerate the computation of expectations of sums of Pauli strings by making use of the Pauli structure vectors representation. The same can be done for gradients of such expectation values using `vectorized_value_and_grad`. Consider a circuit on $n = 3$ qubits, with $k = 4$ Pauli terms and cost function

$$f(w) = \langle Z_0 X_2 \rangle + \langle X_1 Y_2 \rangle + \langle X_0 X_1 Z_2 \rangle + \langle Z_1 \rangle$$

where $\langle Z_0 X_2 \rangle := \langle 0 | U^\dagger(w) Z_0 X_2 U(w) | 0 \rangle$ for some parameterized circuit $U(w)$, and similarly for the other terms. Suppose we have a parameterized circuit defined:

```

1 def param_circuit(params):
2     c = tc.Circuit(n)
3     # circuit details omitted
4     return c

```

then, we can express the expectation of a Pauli term with structure v as:

```

1 def f(w, v):
2     c = param_circuit(w)
3     return tc.templates.measurements.parameterized_measurements(c, v, onehot=True)

```

To batch compute the value and gradient of each term in the cost function, we create a tensor of the Pauli structures:

```

1 structures = tc.array_to_tensor([[3,0,1], #<Z0 X2>
2                                 [0,1,2], #<X1 Y2>
3                                 [1,1,3], #<X0 X1 Z2>
4                                 [0,3,0]]) #<Z1>

```

which we then pass in to the `vectorized_value_and_grad` version of f

```

1 f_vvg = K.vectorized_value_and_grad(f, argnums=0, vectorized_argnums=1)
2 f_vvg(params, structures)

```

Schematically, this returns

$$f_{vvg} \left(w, \begin{pmatrix} \leftarrow v_1 \rightarrow \\ \vdots \\ \leftarrow v_k \rightarrow \end{pmatrix} \right) = \left(\begin{pmatrix} f(w, v_1) \\ \vdots \\ f(w, v_k) \end{pmatrix}, \sum_{i=1}^k \nabla_w f(w, v_i) \right)$$

with each vector v_i corresponding to a different Pauli term.

6.3.4 Batched Machine Learning

In machine learning problems, one often wishes to perform batch computation over (data, label) pairs. This can be done by supplying a tuple of indices to the `vectorized_argnums` argument of `vectorized_value_and_grad`.

Consider the following toy problem where data vectors $x \in [0, 1]^2$ have labels $y \in \{0, 1\}$ and one wishes to train the weights of a parameterized quantum circuit based on a training set of (x, y) pairs. The x vectors are encoded in the angles of an initial set of parameterized gates, with the remaining weights w to be trained to minimize the cost function below:

```

1 def f(x, y, w):
2     c = tc.Circuit(2)
3
4     # encode x in qubit rotations
5     c.rx(0, theta=x[0])
6     c.rx(1, theta=x[1])
7
8     # parameterized circuit to determine label
9     c.rx(0, theta=w[0])
10    c.cnot(0, 1)
11    c.rx(1, theta=w[1])
12
13    yp = c.expectation_ps(z=[1])
14    return K.real(e - y[0]) ** 2, yp

```

Using `vectorized_value_and_grad` now requires the `vectorized_argnums` argument to be a tuple corresponding to the x, y argument indices:

```

1 f_vvg = K.vectorized_value_and_grad(
2     f, argnums=2, vectorized_argnums=(0, 1), has_aux=True
3 )

```

In the $f(x, y, w)$ cost function, x, y are the zeroth and first arguments (which we wish to batch compute over), while w is the second argument, which we wish to take derivatives with respect to. The `has_aux`

argument, if set to True, indicates that the function returns a tuple with only the first element differentiated. In our case, the first output is the loss function to be minimized, and the second auxiliary output is the predicted label \hat{y}_p , which is also helpful to keep for calculating other metrics such as AUC or ROC. The batch computation can then be performed as follows

```
1 X = tc.array_to_tensor([[3, 2], [1, -4], [0, 1]])
2 Y = tc.array_to_tensor([[0], [1], [1]])
3 w = tc.array_to_tensor([0.1, 0.3])
4 f_vvg(X, Y, w)
```

6.3.5 Batched VQE

Consider a cost function defined by a simple parameterized quantum circuit, e.g.,

```
1 def f(w):
2     c = tc.Circuit(2)
3     c.rx(0, theta=w[0])
4     c.cnot(0, 1)
5     c.rx(1, theta=w[1])
6     e = c.expectation_ps(z=[0, 1])
7     return K.real(e)
```

The function value and gradient (with respect to the weights $w[0], w[1]$) can be batch computed for multiple weights simultaneously as follows:

```
1 f_vvag = K.vectorized_value_and_grad(f, argnums=0, vectorized_argnums=0)
2
3 W = tc.array_to_tensor([[0.1, 0.2], [0.3, 0.4]])
4 f_vvag(W)
```

While the above is a toy problem for illustrative purposes, combining this method with `jit` can be a powerful approach to finding ground state energies via VQE, starting from multiple initial points in parameter space and finding the best local minimum reached by gradient descent. The batched VQE workflow, which admits independent optimization loops running simultaneously, is sketched in Figure 8.




Figure 8: Applying `vectorized_value_and_grad` for VQE with batched circuit weights. This enables multiple optimization loop evaluations to be performed at the same time. In the figure, the function f' at the bottom is transformed from the original function f defined at the top via $f' = K.vectorized_value_and_grad(f)$. The weights input is both differentiated and vectorized at the same time.

6.3.6 Batched Monte Carlo trajectory noise simulation

As introduced in Section 5.2.1, by using external random number arrays, we can simulate quantum noise via multiple Monte Carlo trajectories computed in a vectorized parallel fashion. Each trajectory is a specific instance of the noisy circuit, where each quantum channel is replaced by a stochastically chosen operator determined by external random input. The following example shows how to vmap quantum noise (external randomness) in TensorCircuit in a variational quantum algorithm.

```
1 nwires = 6
2
3
4 def f(weights, status):
5     c = tc.Circuit(nwires)
6     # omit the details of constructing a PQC with 'weights' parameters
7     for i in range(nwires):
8         c.depolarizing(i, px=0.2, py=0.2, pz=0.2, status=status[i])
9         # quantum noise controlled by external random number 'status' argument
10        loss = c.expectation_ps(x=[nwires // 2])
11        loss = tc.backend.real(loss)
12        return loss
13
14
15 # get the circuit gradient while vmapping the depolarizing noise
16 f_vg = tc.backend.jit(tc.backend.vvarg(f, argnums=0, vectorized_argnums=1))
17
18 # random number with batch dimension
19 status = tc.backend.implicit_randu(shape=[batch, nwires])
20 f_vg(weights, status)
```

6.4 QuOperator and QuVector

Jupyter notebook: [6-4-quoperator.ipynb](#)

`tc.quantum.QuOperator`, `tc.quantum.QuVector` and `tc.quantum.QuAdjointVector` are data classes which behave like matrices and vectors (columns or rows) when interacting with other ingredients, while their inner structures correspond to tensor networks for efficiency and compactness.

Typical tensor network structures for a `QuOperator`/`QuVector` correspond to Matrix Product Operators (MPO) / Matrix Product States (MPS). The former represents a matrix as:

$$M_{i_1, i_2, \dots, i_n; j_1, j_2, \dots, j_n} = \prod_k T_k^{i_k, j_k},$$

i.e., a product of $d \times d$ matrices $T_k^{i_k, j_k}$, where d is known as the bond dimension. Similarly, an MPS represents a vector as:

$$V_{i_1, \dots, i_n} = \prod_k T_k^{i_k},$$

where the $T_k^{i_k}$ are, again, $d \times d$ matrices. MPS and MPO often occur in computational quantum physics contexts, as they give compact representations for certain types of quantum states and operators. For an introductory review on MPS/MPO in quantum physics, please refer to [16].

`QuOperator`/`QuVector` objects can represent any MPO/MPS, but they can additionally express more flexible tensor network structures. Indeed, any tensor network with two sets of dangling edges of the same dimension (i.e., for each k , the set $\{T_k^{i_k, j_k}\}_{i_k, j_k}$ of matrices has i_k and j_k running over the same index set) can be treated as a `QuOperator`. A general `QuVector` is even more flexible, in that the dangling edge dimensions can be chosen freely; thus, arbitrary tensor products of vectors can be represented.

In this section, we will illustrate the efficiency and compactness of such data structures, and show how they can be integrated seamlessly into quantum circuit simulation tasks. First, consider the following code and tensor diagram (Figure 9) as an introduction to these data class abstractions.

```

1 n1 = tc.gates.Gate(np.ones([2, 2, 2]))
2 n2 = tc.gates.Gate(np.ones([2, 2, 2]))
3 n3 = tc.gates.Gate(np.ones([2, 2]))
4 n1[2] ^ n2[2]
5 n2[1] ^ n3[0]
6
7 matrix = tc.quantum.QuOperator(out_edges=[n1[0], n2[0]], in_edges=[n1[1], n3[1]])
8
9 n4 = tc.gates.Gate(np.ones([2]))
10 n5 = tc.gates.Gate(np.ones([2]))
11
12 vector = tc.quantum.QuVector([n4[0], n5[0]])
13
14 nvector = matrix @ vector # matrix-vector multiplication
15
16 assert type(nvector) == tc.quantum.QuVector
17 nvector.eval_matrix()
18 # array([[16.], [16.], [16.], [16.]])

```

Note that the convention in defining a QuOperator is to first state out_edges (left index or row index of the matrix) and then state in_edges (right index or column index of the matrix). Also note that tc.gates.Gate is just a wrapper for the Node object in the TensorNetwork package.

As seen above, QuOperator/QuVector objects support matrix-matrix or matrix-vector multiplication via the @ operator. Other common matrix/vector operations are also supported:

```

1 matrix.adjoint() # adjoint i.e., conjugate transpose
2 5 * vector # scalar multiplication
3 vector | vector # tensor product
4 matrix.partial_trace([0]) # partial trace (of subsystem 0)

```

Matrix elements of these objects can be extracted via .eval() or .eval_matrix(). The former keeps the shape information of the tensor network while the latter gives the matrix representation (i.e, as a rank 2 tensor).




Figure 9: Tensor network schematic demonstrating the usage of QuOperator and QuVector. Each node is constructed from a tensor and each QuOperator is built by specifying the dangling edges of nodes. From the users' perspective these objects behave as matrices and vectors, with the tensor network engine responsible for maintaining their inner structures and performing calculations.

6.4.1 QuVector as the input state for the circuit

Since a `QuVector` behaves like a regular vector, albeit with a more compact representation, it can be used as the input state to a quantum circuit instead of a state represented as a regular array. The benefit of doing so is memory efficiency. For an n -qubit circuit, regular vector inputs require 2^n complex values to be stored in memory. On the other hand, for an MPS with bond dimension d represented by a `QuVector`, only $O(nd^2)$ complex elements in total need to be stored. Such compact MPS representations can be obtained, for instance, by DMRG [62] calculations, and a DMRG ground state to quantum machine learning model pipeline can thus be built in `TensorCircuit` with the help of this feature.

The following example shows how we input the $|111\rangle$ state, encoded as an MPS, to a quantum circuit. Note how `mps_inputs` argument is used when constructing the circuit.

```
1 n = 3
2 nodes = [tc.gates.Gate(np.array([0.0, 1.0])) for _ in range(n)]
3 mps = tc.quantum.QuVector([nd[0] for nd in nodes])
4 c = tc.Circuit(n, mps_inputs=mps)
5 c.x(0)
6 c.expectation_ps(z=[0]) # 1
```

6.4.2 QuVector as the output state of the circuit

For a given input state, a `tc.Circuit` object can itself be treated as a tensor network with one set of dangling edges corresponding to the output state, and thus the whole circuit object can be regarded as a `QuVector`, obtained via `c.quvector()`. We can then further manipulate the circuit using `QuOperator` objects.

6.4.3 QuOperator as the operator to be measured

As shown in Section 6.2, Hamiltonians can also be represented by a `QuOperator`. This can be a powerful and efficient approach for computing expectation values for some lattice model Hamiltonians like the Heisenberg model or transverse field Ising model (TFIM), as the MPO form of the Hamiltonian for such short-ranged spin models has a very low bond dimension (e.g., $d = 3$ for TFIM). For comparison, for an n -qubit TFIM Hamiltonian, the dense matrix representation stores $O(2^{2n})$ complex elements, the sparse matrix representation stores $O(n2^n)$ complex elements, while the MPO or `QuOperator` representation stores only $18n$ complex elements, which scales linearly with the system size.

Here we show a toy example of measuring the expectation value of an operator represented as a `QuOperator`. The quantity of interest here is $\langle Z_0 Z_1 \rangle$, where the expectation is with respect to the output of a simple circuit which consists of an X gate on one of the qubits. Instead of using the `c.expectation()` API, we use `mpo_expectation`.

```
1 z0, z1 = tc.gates.z(), tc.gates.z()
2 mpo = tc.quantum.QuOperator([z0[0], z1[0]], [z0[1], z1[1]])
3 c = tc.Circuit(2)
4 c.X(0)
5 tc.templates.measurements.mpo_expectation(c, mpo) # -1
```

6.4.4 QuOperator as a quantum gate

Since quantum gates correspond to unitary matrices, an MPO representation for these matrices may allow for significant space efficiencies in some scenarios. A typical example is the multi-controlled gate, which admits a very compact MPO representation that can be characterized by a `QuOperator` in `TensorCircuit`. For an n -qubit gate with $n - 1$ control qubits, the matrix representation for this gate has 2^{2n} elements while the MPO representation can be reduced to bond dimension $d = 2$, leading to only $16n$ elements in memory.

`TensorCircuit` has a built-in way to generate these efficient multi-controlled gates:

```
1 # CCNOT gate
2 c = tc.Circuit(3)
3 c.multicontrol(0, 2, 1, ctrl=[1, 0], unitary=tc.gates.x()) #if q0=1 and q2=0, apply X to q1
```

The 0,2,1 arguments refer to the qubits the gate is applied on (the ordering matters, with the final index referring to the target qubit), the `unitary` argument defines the operation that is applied if all controls are activated and the `ctrl` argument refers to whether the control is activated when the corresponding control qubit is in the 0 state or 1 state. General MPO gates expressed as a `QuOperator` can also be applied via the `c.mpo(*index, mpo=)` API, in the same way that general unitary matrices can be applied via the `c.unitary(*index, unitary=)` API (see Section 3.2).

6.5 Custom contraction settings

Jupyter notebook: [6-5-custom-contraction.ipynb](#)

By default, TensorCircuit uses a greedy tensor contraction path finder provided by the `opt_einsum` package. While this is typically satisfactory for moderately sized quantum circuits, for circuits with 16 qubits or more, we recommend using customized contraction path finders provided by the user or third-party packages.

A simple quantum circuit can be constructed using the following code as a testbed for different contraction methods:

```

1 def testbed():
2     n = 40
3     d = 6
4     param = K.ones([2 * d, n])
5     c = tc.Circuit(n)
6     c = tc.templates.blocks.example_block(c, param, nlayers=d, is_split=True)
7     # the two qubit gate is split and truncated via SVD decomposition
8     return c.expectation_ps(z=[n // 2], reuse=False)
9     # by reuse=False, we compute the expectation as a single tensor network instead of
    first computing the wavefunction

```

By using `tc.templates.blocks.example_block`, a circuit with d layers of $\exp(i\theta ZZ)$ gates and R_x gates is created. When `is_split` is `True`, each two-qubit gate will not be treated as an individual tensor but will be split into two connected tensors via singular value decomposition (SVD), which further simplifies the tensor network structure of the corresponding circuit. The task is to calculate the expectation value of the Z operator on the middle (i.e, $n/2$ -th) qubit.

The API for contraction setup is `tc.set_contractor`. In our example, $2n \times d$ tensors need to be contracted since single-qubit gates can be absorbed into two-qubit gates when `preprocessing` is set to `True` in `set_contractor`. We have some built-in contraction path finder options such as "greedy", "branch", and "optimal" from `opt-einsum` [63], though only the default "greedy" option is suitable for circuit simulation tasks as other options require time exponential in the number of qubits. The `contraction_info` option in this setup API, if set `True`, will print the contraction path information after the path searching. Metrics that measure the quality of a contraction path include

- **FLOPs:** the total number of computational operations required for all matrix multiplications involved when contracting the tensor network via the given path. This metric characterizes the total simulation time.
- **WRITE:** the total size (the number of elements) of all tensors – including intermediate tensors – computed during the contraction.
- **SIZE:** the size of the largest intermediate tensor stored in memory.

Since simulations in TensorCircuit are AD-enabled, where all intermediate results need to be cached and traced, the more relevant spatial cost metric is write instead of size.

6.5.1 Customized contraction path finder

For large quantum circuits, the performance of the default "greedy" contraction path finder may not be satisfactory. If this is the case, a custom contraction path finder can be used to enhance the performance of contraction by finding better paths in terms of flops (time) and writes (space). Here we use the path finder provided by the third-party `cotengra` package, a python library for contracting tensor networks or computing einsum expressions. The way to use the `cotengra` path finder in TensorCircuit is as follows:

```

1 import cotengra as ctg
2
3 opt = ctg.ReusableHyperOptimizer(
4     methods=["greedy", "kahypar"],
5     parallel=True,
6     minimize="write",
7     max_time=120,
8     max_repeats=1024,
9     progbar=True,
10 )
11 tc.set_contractor("custom", optimizer=opt, preprocessing=True, contraction_info=True)
12 testbed()

```

A number of parameters are used to configure a contraction path finder `opt` in cotengra: `method` decides the strategy this path finder will be based on. `minimize` decides the score function you want to minimize during the path finding, and can be set as "write", "flops", "size" or a combination of these. A time limit and a limit on the number of trial contraction trees can also be set using `max_time` and `max_repeats` respectively. For more details, refer to the cotengra documentation. You can also design your own contraction path finder as long as you provide an `opt` function compatible with the interface of the `opt_einsum` optimizer.

6.5.2 Subtree reconfiguration

Given a contraction path, e.g. given by a "greedy" search, its performance can be further enhanced by conducting a so-called subtree reconfiguration. This process repeatedly optimizes subtrees of the whole contraction tree, and in practice often results in a better contraction path. This can be done in TensorCircuit as follows:

```

1 opt = ctg.ReusableHyperOptimizer(
2     minimize="combo",
3     max_repeats=1024,
4     max_time=120,
5     progbar=True,
6 )
7
8
9 def opt_reconf(inputs, output, size, **kws):
10     tree = opt.search(inputs, output, size)
11     tree_r = tree.subtree_reconfigure_forest(
12         progbar=True, num_trees=10, num_restarts=20, subtree_weight_what=("size",)
13     )
14     return tree_r.get_path()
15
16
17 tc.set_contractor(
18     "custom",
19     optimizer=opt_reconf,
20     contraction_info=True,
21     preprocessing=True,
22 )
23
24
25 testbed()

```

Notice that `subtree_reconfigure_forest` is used after finding a contraction tree. In this function, you can set the number of trees in the random forest whose contraction paths will be updated, and also the metric to be optimized in the subtrees. In the above example, a user customized function `opt_reconf` is fed into contractor setup as a legal contraction path finder.

As mentioned earlier, there are three metrics to measure the quality of a contraction path. By setting different score functions (changing the `minimize` parameter), the resulting contraction path given by these contractors will exhibit different properties. Table 5 summarizes the contraction performance of different contraction strategies for our example case ($n = 40, d = 6$). As we can see, the cotengra optimizer and subtree reconfiguration can greatly improve the quality of the contraction path and improve the efficiency of quantum circuit simulation. For example, we gain more than a factor of two improvement in

contractor	reconfiguration	log10[FLOPs]	log2[SIZE]	log2[WRITE]
default		7.373	12	20.171
cotengra("flops")		7.080	13	20.493
cotengra("flops")	subtree("flops")	7.006	12	20.069
cotengra("flops")	subtree("size")	7.006	12	20.075
cotengra("write")		7.442	14	19.061
cotengra("write")	subtree("flops")	7.000	12	19.988
cotengra("write")	subtree("size")	7.017	12	19.958
cotengra("combo")		7.480	14	19.061
cotengra("combo")	subtree("flops")	7.003	12	20.000
cotengra("combo")	subtree("size")	7.011	12	19.885

Table 5: Performance of different contractor settings which include the default opt_einsum contractor and cotengra contractors with and without subtree reconfiguration. Parameters in the brackets indicate the score function used during the path searching and reconfiguration. For cotengra contractors, we set max_repeats=1024, max_time=120 and method=["greedy", "kahypar"]. "combo" means the score function is a combination of "flops" and "write" (flops+64×write by default). For subtree reconfiguration, we set num_trees=20, num_restarts=20. Results shown are from one run, and performance may vary from run to run since these algorithms are intrinsically random.

simulation time and simulation space compared to the default contractor, and the degree of improvement can increase for larger system sizes.

6.6 Advanced automatic differentiation

Jupyter notebook: [6-6-advanced-automatic-differentiation.ipynb](#)

TensorCircuit provides backend-agnostic wrappers to a number of advanced AD features, useful in a variety of quantum circuit simulation scenarios. In the remainder of this section we will illustrate these using the following circuit example:

```

1 n = 6
2 nlayers = 3
3
4
5 def ansatz(thetas):
6     c = tc.Circuit(n)
7     for j in range(nlayers):
8         for i in range(n):
9             c.rx(i, theta=thetas[j])
10            for i in range(n - 1):
11                c.cnot(i, i + 1)
12            return c
13
14
15 def psi(thetas):
16     c = ansatz(thetas)
17     return c.state()

```

Jacobian (jacfwd and jacrev). Given an n -input, m -output function f the $n \times m$ Jacobian matrix is given by

$$J_f := \frac{\partial f}{\partial x} = \begin{pmatrix} \frac{\partial f_1}{\partial x_1} & \cdots & \frac{\partial f_1}{\partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_m}{\partial x_1} & \cdots & \frac{\partial f_m}{\partial x_n} \end{pmatrix}$$

By the chain rule, the Jacobian matrix of a composition of functions is the product of the Jacobians of the composed functions (evaluated at appropriate points). e.g. if $h : \mathbb{R}^n \rightarrow \mathbb{R}^p$, $g : \mathbb{R}^p \rightarrow \mathbb{R}^q$, $f : \mathbb{R}^q \rightarrow \mathbb{R}^m$

and $y : \mathbb{R}^n \rightarrow \mathbb{R}^m$ with $y(x) = f(g(h(x)))$ then

$$\begin{aligned}\frac{\partial y}{\partial x} &= \frac{\partial f(b)}{\partial b} \cdot \frac{\partial g(a)}{\partial a} \cdot \frac{\partial h(x)}{\partial x} \\ &= J_f(b) \cdot J_g(a) \cdot J_h(x)\end{aligned}$$

where $a = h(x), b = g(a)$ and \cdot denotes matrix multiplication. Forward mode AD and reverse mode AD ('backpropagation') are two approaches to computing a composite Jacobian, and differ in the order in which the products are computed. Forward mode AD computes the above product from right to left, i.e. $J_y = J_f(b) \cdot (J_g(a) \cdot J_h(x))$ at a cost of $pqn + qnm$ multiplications. Taking f to be the output state $\psi(\theta)$ of the above circuit, this is computed as

```
1 thetas = K.implicit_randn([nlayers])
2 jac_fwd_function = K.jacfw(psi)
3 jac_fw = jac_fwd_function(thetas)
```

Reverse mode AD computes the product from left to right, i.e. $J_y = (J_f(b) \cdot) J_g(a) \cdot J_h(x)$ at a cost of $mpq + pnm$ multiplications:

```
1 jac_rev_function = K.jacrev(psi)
2 jac_rev = jac_rev_function(thetas)
```

The relative efficiency of these methods depends on the input and output dimensions of the functions involved. For instance, when $p = q$ forward mode AD is advantageous if $n \ll m$ (i.e., the input dimension is much smaller than the output dimension, corresponding to a 'tall' Jacobian) and vice versa for reverse mode AD.

Jacobian-vector product (jvp). Computing the product of the Jacobian with a vector v (i.e. the directional derivative) can be a useful primitive as it utilizes forward mode AD and is suitable when the output dimension is much larger than the input. For instance, setting $v = e_i$ (i.e. the vector with a 1 in the i -th coordinate and zeroes elsewhere) gives the vector of partial derivatives

$$J_f e_i = \left(\frac{\partial f_1}{\partial x_i}, \dots, \frac{\partial f_m}{\partial x_i} \right)^\top$$

Taking $v = (1.0, 0, 0)$, the value of ψ and the Jacobian-vector product $\frac{\partial \psi}{\partial \theta_0}$ can be evaluated (e.g., at the point $\theta = (0.1, 0.2, 0.3)$) as follows:

```
1 state, partial_psi_partial_theta0 = K.jvp(
2     psi,
3     tc.array_to_tensor([0.1, 0.2, 0.3]),
4     tc.array_to_tensor([1.0, 0, 0], dtype="float32"),
5 )
```

Quantum Fisher Information (qng). The Quantum Fisher Information (QFI) is an important concept in quantum information, and can be utilized in so-called quantum natural gradient descent optimization [64] as well as variational quantum dynamics [65, 66].

There are several variants of QFI-like quantities, all of which depend on the evaluation of terms of the form $\langle \partial_i \psi | \partial_j \psi \rangle - \langle \partial_i \psi | \psi \rangle \langle \psi | \partial_j \psi \rangle$. Such quantities are easily obtained with advanced AD frameworks, by first computing the Jacobian for the output state and then vmapping the inner product over Jacobian rows. The detailed efficient implementation can be found at the codebase. Here we directly call the corresponding API to obtain the quantum natural gradient.

```
1 from tensorcircuit.experimental import qng
2
3 # function to get qfi given circuit parameters
4 qfi_fun = K.jit(qng(psi))
5
6 # suppose the vanilla circuit gradient is 'grad'
7 # then we can obtain quantum natural gradient as
8 ngrad = tc.backend.solve(qfi_fun(thetas), grad, assume_a="sym")
```

Hessian (hessian). The Hessian $H_{ij} = \frac{\partial \langle H \rangle_\theta}{\partial \theta_i \partial \theta_j}$ of a parameterized quantum circuit can be computed as (taking $H = Z_0$ for simplicity)

```

1 def h(thetas):
2     c = ansatz(thetas)
3     return c.expectation_ps(z=[0])
4
5 # hess is the Hessian function which takes thetas as input
6 hess = K.hessian(h)

```

which can then also be jitted to make multiple evaluations more efficient:

```

1 hess_jit = K.jit(hess)

```

Information on the Hessian matrix may be useful in investigating loss landscapes, or for second order optimization routines.

$\langle \psi | \mathbf{H} | \partial \psi \rangle$. In variational quantum dynamics problems (see, e.g. [65]) one often wishes to compute quantities of the form

$$\langle \psi(\theta) | H \frac{\partial |\psi(\theta)\rangle}{\partial \theta_i}.$$

This can be done via the `stop_gradient` API, which prevents certain parameters from being differentiated. Again, taking $H = Z_0$, we can define an appropriate function for which only the parameters corresponding to $|\psi(\theta)\rangle$ will be differentiated:

```

1 z0 = tc.quantum.PauliStringSum2Dense([[3, 0, 0, 0, 0, 0]])
2
3
4 def h(thetas):
5     w = psi(thetas)
6     w_left = K.conj(w)
7     w_right = K.stop_gradient(w)
8     w_left = K.reshape([1, -1])
9     w_right = K.reshape([-1, 1])
10    e = w_left @ z0 @ w_right
11    return K.real(e)[0, 0]

```

Then, gradients can be computed as usual:

```

1 psi_h_partial_psi = K.grad(h)

```

With the advanced automatic differentiation infrastructure, we can obtain quantum circuit gradient related quantities, such as those listed above, much more quickly than via traditional quantum software that utilizes parameter shifts to evaluate gradients. In Figure 10, we show the acceleration of QFI and Hessian computations compared to Qiskit. The benchmark code is detailed in [examples/gradient_benchmark.py](#). From the data, we see that for even moderate-sized quantum circuits, TensorCircuit can achieve a speedup over Qiskit of nearly a *million times*.

7 Integrated examples

7.1 Molecular VQE with TensorCircuit and OpenFermion

Jupyter notebook: [tutorials/vqe_h2o.ipynb](#)

Background. Quantum computing is envisioned to be a powerful tool for quantum simulation and quantum chemistry tasks [67, 68]. In the *highlighted features* below we show how to interface TensorCircuit with OpenFermion [69] to compute the ground state energy of an H_2O molecule, while in the *benchmarking* section we compare the performance of TensorCircuit with other software for the VQE energy computation of a transverse field Ising model.

Highlighted features. OpenFermion is an open-source Python package that provides an efficient interface between quantum chemistry and quantum computing. In particular, it provides a convenient method for generating molecular Hamiltonians and converting them into qubit Hamiltonians compatible with




Figure 10: Acceleration factor for TensorCircuit over Qiskit for the evaluation of QFI and Hessian information. For 4 and 6 qubit systems, we use PQC comprising two blocks of CNOT+Rx+Rz+Rx gates, while for the larger systems we use four such blocks. The simulation runs on AMD EPYC 7K62 CPU 2.60GHz, and TensorCircuit results use the JAX backend. As the Qiskit running times are long (e.g., the 12-qubit QFI calculation takes more than 20,000s \sim 5.5 hours) the acceleration factors presented above are based on only a single Qiskit evaluation. In contrast, the TensorCircuit running times are low enough (0.026s for 12-qubit QFI) that we average over multiple runs.

simulating in quantum circuits. To make use of this, we provide the `tc.templates.chems.get_ps` API, which provides an interface between TensorCircuit and OpenFermion, and converts the OpenFermion qubit Hamiltonian object into the Pauli structures and weights tensors used by TensorCircuit (see Section 6.2.1). The relevant code snippet used to generate the Hamiltonian representation in TensorCircuit via OpenFermion is shown below.

```

1 from openfermion.chem import MolecularData, geometry_from_pubchem
2 from openfermion.transforms import get_fermion_operator, jordan_wigner
3 from openfermionpyscf import run_pyscf
4
5 multiplicity = 1
6 basis = "sto-3g"
7 # 14 spin orbitals for H2O
8 geometry = geometry_from_pubchem("h2o")
9 molecule = MolecularData(geometry, basis, multiplicity)
10 # obtain H2O molecule object
11 molecule = run_pyscf(molecule, run_mp2=True, run_cisd=True, run_ccsd=True, run_fci=True)
12 print(molecule.fci_energy, molecule.ccsd_energy, molecule.hf_energy)
13
14 mh = molecule.get_molecular_hamiltonian()
15 # get fermionic Hamiltonian
16 fh = get_fermion_operator(mh)
17 # get qubit Hamiltonian via Jordan-Wigner transformation
18 jw = jordan_wigner(fh)
19 # converting to Pauli structures in tc
20 structures, weights = tc.templates.chems.get_ps(jw, 14)
21 # build sparse numpy matrix representation for the Hamiltonian
22 ma = tc.quantum.PauliStringSum2C00_numpy(structures, weights)

```

Benchmarking. We provide some benchmark data for TFIM VQE evaluation here. Specifically, we compute the TFIM energy expectation value and corresponding circuit gradients (with respect to the circuit parameters). With TensorCircuit, we use the (potentially less efficient) explicit for loop to perform the Pauli string summation (see Section 6.2.2) to obtain a fair comparison, with our benchmarking focus on the efficiency of evaluating the PQC and its gradients. See [benchmarks](#) for benchmark setup and details. Results are summarized in Table 6. Note that while we consider the TFIM model in our benchmark tests, this only differs from molecular systems in the number and type of Pauli strings that appear in the Hamiltonian.

We do not perform similar benchmark tests with other common quantum simulator packages such as Qiskit, Cirq and ProjectQ, as these do not support automatic differentiation, relying instead on the parameter shift technique to evaluate gradients. As the computational complexity of this approach scales linearly with the number of variational parameters, this is inefficient for circuits with a large number of parameters. For example, using Qiskit with its built-in parameter-shift gradient framework, the task in

PQC	$n = 10, d = 3$	$n = 16, d = 16$	$n = 22, d = 11$
PennyLane (CPU)	0.012	0.31	24.84
PennyLane (GPU)	0.067	0.68	OOM
TensorFlow Quantum	0.005	0.026	0.68
Qibo (GPU)	0.033	0.198	OOM
TensorCircuit (CPU)	0.00077	0.078	4.70
TensorCircuit (GPU)	0.0026	0.023	0.19

Table 6: Performance benchmarks (running time in seconds) for value and gradient evaluations of n -qubit, d -layer parameterized quantum circuits with a one-dimensional open boundary condition TFIM Hamiltonian objective function. Each layer of the hardware efficient ansatz comprises $n - 1$ Rzz gates arranged in a cascading ladder layout and one layer of Rx gates. TensorCircuit results use the JAX backend. GPU simulations use the Nvidia V100 32G GPU while CPU simulations use Intel(R) Xeon(R) Platinum 8255C CPU @ 2.50GHz. OOM indicates that the GPU memory is insufficient to run the corresponding benchmark code. Note that TensorFlow Quantum only currently supports quantum circuit simulations on CPU. Qibo results use the TensorFlow backend since this is the only Qibo backend that supports automatic differentiation. PennyLane CPU results use the PennyLane-Lightning C++ backend, while the GPU results make use of vectorized parallelism and JIT via the JAX backend. At the time of writing, the GPU version of PennyLane-Lightning does not support automatic differentiation.

Table 6 (with $n = 10, d = 3$ takes 125s per iteration – more than 10^5 times slower than TensorCircuit). Similar results have already been presented in Section 6.6.

7.2 Quantum machine learning

Jupyter notebook: [tutorials/mnist_qml.ipynb](#)

Background. Quantum and hybrid quantum-classical neural networks are popular approaches to NISQ era quantum computing, and both can be easily modelled and tested in TensorCircuit. In the *highlighted features* below, we illustrate how to build a hybrid machine learning pipeline in TensorCircuit, and in the *benchmarking* section we compare TensorCircuit with other quantum software for performing batched supervised learning on the MNIST dataset using a parameterized quantum circuit.

Highlighted features. Seamless integration of quantum and classical neural networks can be obtained by wrapping the TensorCircuit `tc.Circuit` object (with weights as input and expectation value as output) in a `QuantumLayer`, which is a subclass of the Keras `Layer`. The following code snippet shows how such a wrapper is implemented and used:

```

1  def qml_ys(x, weights, nlayers):
2      n = 9
3      weights = tc.backend.cast(weights, "complex128")
4      x = tc.backend.cast(x, "complex128")
5      c = tc.Circuit(n)
6      for i in range(n):
7          c.rx(i, theta=x[i])
8      for j in range(nlayers):
9          for i in range(n - 1):
10             c.cnot(i, i + 1)
11             for i in range(n):
12                 c.rx(i, theta=weights[2 * j, i])
13                 c.ry(i, theta=weights[2 * j + 1, i])
14             ypreds = []
15             for i in range(n):
16                 ypred = c.expectation([tc.gates.z(), (i,)])
17                 ypred = tc.backend.real(ypred)
18                 ypred = (tc.backend.real(ypred) + 1) / 2.0
19                 ypreds.append(ypred)
20             # return <z_i> as an n dimensional vector
21             return tc.backend.stack(ypreds)
22
23

```

batch size	32	128	512
PennyLane (CPU)	0.58	2.21	8.80
PennyLane (GPU)	0.042*	0.0089	0.020
TensorFlow Quantum	0.058	0.24	0.49
TensorCircuit (CPU)	0.0070	0.021	0.085
TensorCircuit (GPU)	0.0035	0.0039	0.0054

Table 7: Performance benchmarks (running time in seconds) for value and gradient evaluation of a PQC ($n = 10$ qubits, circuit depth $p = 3$) with classification square distance error as the objective function and batched dataset input. TensorCircuit results use the JAX backend. GPU simulations use the Nvidia V100 32G GPU while CPU simulations use Intel(R) Xeon(R) Platinum 8255C CPU @ 2.50GHz. Note that TensorFlow Quantum only runs circuit simulations on CPU. PennyLane CPU results use the PennyLane-Lightning C++ backend, together with the PennyLane native batch (broadcast) paradigm. PennyLane GPU results make use of vectorized parallelism and JIT via the JAX backend. *The PennyLane running time on GPU is indeed higher for smaller batch size.

```

24 # wrap the quantum function in a Keras layer
25 ql = tc.keras.QuantumLayer(partial(qml_ys, nlayers=nlayers), [(2 * nlayers, 9)])
26 # build the hybrid Keras model with quantum and classical parts
27 model = tf.keras.Sequential([ql, tf.keras.layers.Dense(1, activation="sigmoid")])
28
29 # train as a normal Keras model
30 model.compile(
31     loss=tf.keras.losses.BinaryCrossentropy(),
32     optimizer=tf.keras.optimizers.Adam(0.01),
33     metrics=[tf.keras.metrics.BinaryAccuracy()],
34 )
35 model.fit(x_train, y_train, batch_size=32, epochs=100)

```

Benchmarking. We benchmark TensorCircuit against other software for binary classification (‘3’ vs. ‘6’) of the MNIST dataset, using quantum machine learning with batched inputs. For software with only parameter-shift gradient support, each PQC must be evaluated $O(np)$ times, where np is the number of circuit parameters. This is much slower than AD-enabled simulators, where only one evaluation of the PQC suffices to obtain all circuit weight gradients. Therefore, we only compare TensorCircuit with other AD enabled software such as TensorFlow Quantum and PennyLane (With PennyLane, we utilize its JAX backend simulator, and use jit and vmap tricks to increase performance). The TensorCircuit results use the default greedy contraction path finder, and further improvements are possible with customized contraction path finders. See [benchmarks](#) for full benchmark details and Table 7 for results.

From the benchmarks on the standard VQE and QML task sets in the previous two examples, we see that TensorCircuit can indeed bring substantial speedups in quantum circuit simulation. The acceleration is more impressive on GPU, especially when the circuit size or the batch dimension is large.

7.3 Demonstration of barren plateaus

Jupyter notebook: [tutorials/barren_plateaus.ipynb](#)

Background. The so-called barren plateaus phenomenon refers to gradients of random circuits vanishing exponentially quickly as the number of qubits or the circuit depth increases [13]. To demonstrate this numerically requires computing the circuit gradient variance over different circuit structures (i.e., choice of random gates in the circuit) and circuit weights. Gradients can be obtained via automatic differentiation, while the different circuit weights can be vectorized and jitted to boost performance. In addition, as in this example, the different circuit structures can also be vectorized and jitted to obtain further speed up.

Highlighted features. This example showcases how jit and vmap can be applied to different circuit architectures. The ability to vmap circuit structures was introduced in Section 6.3.2. Here, we give more details on how to encode different circuit structures via a tensor parameter as input. The core part of the circuit construction is as follows.

	PennyLane (CPU)	TensorFlow Quantum	TensorCircuit (CPU)	TensorCircuit (GPU)
time (s)	155.28	6.24	0.12	0.011

Table 8: Performance benchmark for gradient variance evaluation over different random circuit architectures and circuit weights (10 qubits, 10 layers, 100 different circuits). TensorCircuit results use the JAX backend. PennyLane results use the PennyLane-Lightning backend with highly optimized C++ code. GPU simulations use the Nvidia V100 32G GPU while CPU simulations use Intel(R) Xeon(R) Gold 6133 CPU @ 2.50GHz. Note that TensorFlow Quantum only supports CPU for quantum circuit simulation.

```

1 Rx = tc.gates.rx
2 Ry = tc.gates.ry
3 Rz = tc.gates.rz
4
5 # params is a tensor for the circuit weights with shape [n_qubits, n_layers]
6 # seeds is a tensor for the circuit structure with shape [n_qubits, n_layers]
7
8 c = tc.Circuit(n_qubits)
9 for l in range(n_layers):
10     for i in range(n_qubits):
11         c.unitary_kraus(
12             [Rx(params[i, l]), Ry(params[i, l]), Rz(params[i, l])],
13             i,
14             prob=[1 / 3, 1 / 3, 1 / 3],
15             status=seeds[i, l],
16         )
17     for i in range(n_qubits - 1):
18         c.cz(i, i + 1)

```

The `seeds` tensor controls the circuit architecture via the `unitary_kraus` API. This API tells the circuit to stochastically attach one gate from Rx, Ry, Rz with probability 1/3 each. The `status` argument externalizes the random number generation (see Section 5.2.1). Namely, when `status` is less than 1/3, the first gate is applied, when `status` is in the range $[1/3, 2/3]$, the second gate in the list is applied and so on. Therefore, by generating a random number array `seeds = K.implicit_randu(size=[n_qubits, n_layers])`, we can generate different random circuit s. The simulation over different architectures can be vectorized and jitted by generating `seeds` with an extra batch dimension corresponding to the number of different architectures you wish to compute in parallel.

Benchmarking. We benchmark the gradient variance computation over different random circuit weights and different circuit structures using TensorFlow Quantum, PennyLane and TensorCircuit, (see details in

[examples/bp_benchmark.py](#)). The gradient computational times for 100 random circuit constructions, each a 10-qubit, 10-layer circuit, are shown in Table 8. With 10×10 circuit weights and 100 different circuits, the combination of `vmap` and `jit` provides TensorCircuit with a more than *five hundred times* speedup over TensorFlow Quantum for this task. As PennyLane lacks the capability to batch compute the trainable parameters and different circuit structures, the gradient evaluations must be performed sequentially by a naïve loop, which leads to significantly longer times even on the fastest PennyLane-Lightning backend.

The choice of benchmark problem sets: The previous three subsections gave benchmark results for VQE, QML and the barren plateaus problems, respectively. These three tasks were chosen to showcase and benchmark the paradigms and features that quantum software can support. VQE problems can benefit from AD, QML problems can benefit from AD and VMAP on non-differentiable inputs, while the barren plateaus problem can be accelerated by AD as well as VMAP on differentiable inputs and on circuit structures. In all cases, JIT and GPU support (if compatible with AD and VMAP) can further increase the efficiency of the corresponding quantum simulations.

7.4 Very large circuit simulation

Python script: [examples/vqe_extra_mpo.py](#)

Background. As previously mentioned, tensor network quantum simulators do not face the same memory bottlenecks that limit full state simulators, and can thus simulate larger numbers of qubits as long as circuit connectivity and depth are reasonably low. In the *highlighted features* and *benchmarking* sections below, we consider a one dimensional TFIM VQE workflow on 600 qubits with seven layers of two-qubit gates, arranged in a cascading ladder layout, that estimates energy with more than 99% accuracy compared to ground truth. The Hamiltonian we consider has open boundary condition and sits at the critical point, which is believed to be the most difficult point to simulate on the phase diagram.

Highlighted features. The combination of the MPO formalism with the cotengra path finder allows us to simulate circuits with very large qubit counts. Specifically, we utilize (i) the advanced cotengra path finder equipped with subtree reconfiguration post-processing, and with a setup similar to Section 6.5.2; (ii) the space efficient MPO representation to evaluate the quantum expectation of the TFIM Hamiltonian (see Section 6.2.3). For the circuit construction, the `split` configuration was used to decompose the parameterized ZZ gates, using SVD to reduce the bond dimension to 2, which further simplifies the tensor network structure to be contracted (see Figure 11). Such a two-qubit gate decomposition is implemented as follows.

```
1 split_conf = {
2     "max_singular_values": 2,
3     "fixed_choice": 1,
4 }
5 # set the SVD decomposition option in circuit level
6 c = tc.Circuit(n, split=split_conf)
7 # or set the SVD decomposition option in gate level
8 c.exp1(i, i + 1, theta=param, unitary=tc.gates._xx_matrix, split=split_conf)
```

The ground energy for reference is obtained via two-site DMRG using the Quimb package.




Figure 11: SVD decomposition of a two-qubit gate as a tensor network. This operation can be used to reduce the complexity of the underlying network. While the bond dimension d after decomposition for general two-qubit gates is 4, for certain types of two-qubit gates it can be lower. For example, in our case, where two-qubit parameterized gates are of the form $\exp(i\theta XX)$, $d = 2$.

Benchmarking. The parameterized circuit we utilized has n qubits and 7 layers, and contains $21n$ single-qubit gates and $7n$ two-qubit gates. We note that our circuit ansatz relies on a succinct tensor network representation for its efficient simulation. In each layer, $n - 1$ two-qubit gates are arranged in a cascading ladder layout, i.e. the gates are applied on the $(0, 1), (1, 2), (2, 3) \dots (n - 2, n - 1)$ qubit pairs. This ladder layout can spread quantum information and entanglement across the entire chain in only one layer of two-qubit gates, and the resulting state can not readily be simulated by, for instance, taking advantage of a restricted causal light cone as in [70].

Benchmark results for $n = 200, 400, 600$ are summarized in Table 9, with times per computational step corresponding to the evaluation of both the energy expectation value and the circuit gradients. We find that high-accuracy results are obtained even for large systems without any fine-tuning of the optimizer. This may be explained by the fact that the 1D TFIM has a simple ground state and entanglement structure, and overparametrization (the model has a large number of independent trainable parameters) may also contribute to the accuracy obtained [71].

number of qubits	$\log_2[\text{WRITE}]$	time for one step	energy accuracy reached
200	27.3	5.7	99.6%
400	29.3	11.8	99.5%
600	31.5	18.2	99.4%

Table 9: Performance benchmarks (running time in seconds) for a large scale VQE task with different qubit counts. TensorCircuit results use the TensorFlow backend and run on Nvidia A100 40G GPU. Note that the VQE optimization hyperparameters were not tuned, so the energy accuracy obtained only represents a lower bound on the performance of the current method. Time for one step includes computing both the energy expectation value and the circuit gradients. The full optimization requires thousands of iterations, with the Adam gradient descent optimizer used initially, followed by an SGD optimizer in the later stages of the process.

8 Outlook and concluding remarks

We have introduced TensorCircuit, an open-source Python package, designed to meet the requirements of larger and more complex quantum computing simulations. TensorCircuit is built on top of, and incorporates, all the main engineering paradigms from modern machine learning libraries, and its flexible and customizable tensor network engine enables high-performance circuit computation.

Outlook. We will continue the development of TensorCircuit, towards delivering a more efficient and elegant, full-featured and ML-compatible quantum software package. At the top of our priority list are:

1. Better tensor network contraction path finders: integrate more advanced algorithms and machine learning techniques for optimal contraction path searching.
2. Pulse level optimization and quantum control: enable end-to-end differentiable pulse level optimization and optimal quantum control schedules [72, 73].
3. Distributed quantum circuit simulation: enable tensor network parallel slicing and distributed computation on multiple hosts.
4. Approximate circuit simulation based on MPS: introduce TEBD-like algorithms [74, 75] to approximately simulate quantum circuits with large size and depth.
5. More quantum-aware or manifold-aware optimizers: include optimizers such as SPSA [76], rotate [77], and Riemannian optimizers [78].
6. Quantum applications: develop application-level libraries for quantum computing for finance, materials, energy, biology, drug discovery, climate prediction and more.

With the continued rapid progress in quantum computing theory and hardware, our hope is that TensorCircuit, a quantum simulator platform designed for the NISQ era, will play an important role in the academic and commercial progress of this exciting field.

Acknowledgements: The authors would like to thank our teammates at the Tencent Quantum Laboratory for supporting this project, and Tencent Cloud for providing computing resources. Shi-Xin Zhang would like to thank Rong-Jun Feng, Sai-Nan Huai, Dan-Yu Li, Zi-Xiang Li, Lei Wang, Hao Xie, Shuai Yin and Hao-Kai Zhang for their helpful discussions.

References

- [1] Michael A. Nielsen and Isaac L. Chuang. “Quantum computation and quantum information: 10th anniversary edition”. Cambridge University Press. USA (2011). 10th edition.
- [2] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. “Tensorflow: A system for large-scale machine learning”. In 12th USENIX symposium on operating systems design and implementation (OSDI 16). Pages 265–283. (2016).

- [3] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. “JAX: composable transformations of Python+NumPy programs”. <http://github.com/google/jax> (2018).
- [4] 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. “Pytorch: An imperative style, high-performance deep learning library”. In *Advances in Neural Information Processing Systems*. Volume 32. (2019).
- [5] Johnnie Gray. “cotengra”. <https://github.com/jcmgray/cotengra> (2020).
- [6] Johnnie Gray and Stefanos Kourtis. “Hyper-optimized tensor network contraction”. *Quantum* **5**, 410 (2021).
- [7] John Preskill. “Quantum computing in the nisq era and beyond”. *Quantum* **2**, 79 (2018).
- [8] Kishor Bharti, Alba Cervera-Lierta, Thi Ha Kyaw, Tobias Haug, Sumner Alperin-Lea, Abhinav Anand, Matthias Degroote, Hermann Heimonen, Jakob S. Kottmann, Tim Menke, Wai-Keong Mok, Sukin Sim, Leong-Chuan Kwek, and Alán Aspuru-Guzik. “Noisy intermediate-scale quantum algorithms”. *Reviews of Modern Physics* **94**, 015004 (2022).
- [9] M. Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C. Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R. McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, and Patrick J. Coles. “Variational quantum algorithms”. *Nature Reviews Physics* **3**, 625–644 (2021).
- [10] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O’Brien. “A variational eigenvalue solver on a photonic quantum processor”. *Nature communications* **5**, 4213 (2014).
- [11] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. “A quantum approximate optimization algorithm”. *arXiv eprint* (2014) arXiv:1411.4028.
- [12] Abhinav Kandala, Antonio Mezzacapo, Kristan Temme, Maika Takita, Markus Brink, Jerry M Chow, and Jay M Gambetta. “Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets”. *Nature* **549**, 242–246 (2017).
- [13] Jarrod R. McClean, Sergio Boixo, Vadim N Smelyanskiy, Ryan Babbush, and Hartmut Neven. “Barren plateaus in quantum neural network training landscapes”. *Nature Communications* **9**, 4812 (2018).
- [14] Eric R. Anschuetz. “Critical points in quantum generative models”. *arXiv eprint* (2021) arXiv:2109.06957.
- [15] M-H Yung, Jorge Casanova, Antonio Mezzacapo, Jarrod Mcclean, Lucas Lamata, Alan Aspuru-Guzik, and Enrique Solano. “From transistor to trapped-ion computers for quantum chemistry”. *Scientific reports* **4**, 3589 (2014).
- [16] Ulrich Schollwöck. “The density-matrix renormalization group in the age of matrix product states”. *Annals of Physics* **326**, 96–192 (2011).
- [17] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. “Deep learning”. *Nature* **521**, 436–444 (2015).
- [18] Michael Bartholomew-Biggs, Steven Brown, Bruce Christianson, and Laurence Dixon. “Automatic differentiation of algorithms”. *J. Comput. Appl. Math.* **124**, 171–190 (2000).
- [19] Atılım Güneş Baydin, Barak A. Pearlmutter, Alexey Andreyevich Radul, Jeffrey Mark Siskind, Atılım Gunes Baydin, Barak A. Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. “Automatic differentiation in machine learning: A survey”. *J. Mach. Learn. Res.* **18**, 1–43 (2018).

- [20] Jun Li, Xiaodong Yang, Xinhua Peng, and Chang-Pu Sun. “Hybrid quantum-classical approach to quantum optimal control”. *Phys. Rev. Lett.* **118**, 150503 (2017).
- [21] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. “Evaluating analytic gradients on quantum hardware”. *Phys. Rev. A* **99**, 032331 (2019).
- [22] MD SAJID ANIS et al. “Qiskit: An open-source framework for quantum computing”. <https://github.com/qiskit> (2021).
- [23] Cirq Developers. “Cirq”. <https://github.com/quantumlib/Cirq> (2021).
- [24] Damian S. Steiger, Thomas Häner, and Matthias Troyer. “ProjectQ: an open source software framework for quantum computing”. *Quantum* **2**, 49 (2018).
- [25] Huawei HiQ team. “Huawei HiQ: A high-performance quantum computing simulator and programming framework”. <http://hiq.huaweicloud.com>.
- [26] Krysta Svore, Alan Geller, Matthias Troyer, John Azariah, Christopher Granade, Bettina Heim, Vadym Kliuchnikov, Mariia Mykhailova, Andres Paz, and Martin Roetteler. “Q# enabling scalable quantum computing and development with a high-level dsl”. In *Proceedings of the real world domain specific languages workshop 2018*. Pages 1–10. (2018).
- [27] Stavros Efthymiou, Sergi Ramos-Calderer, Carlos Bravo-Prieto, Adrián Pérez-Salinas, Diego García-Martín, Artur Garcia-Saez, José Ignacio Latorre, and Stefano Carrazza. “Qibo: a framework for quantum simulation with hardware acceleration”. *Quantum Science and Technology* **7**, 015018 (2022).
- [28] Yasunari Suzuki, Yoshiaki Kawase, Yuya Masumura, Yuria Hiraga, Masahiro Nakadai, Jiabao Chen, Ken M. Nakanishi, Kosuke Mitarai, Ryosuke Imai, Shiro Tamiya, Takahiro Yamamoto, Tennin Yan, Toru Kawakubo, Yuya O. Nakagawa, Yohei Ibe, Youyuan Zhang, Hirotsugu Yamashita, Hikaru Yoshimura, Akihiro Hayashi, and Keisuke Fujii. “Qulacs: a fast and versatile quantum circuit simulator for research purpose”. *Quantum* **5**, 559 (2021).
- [29] Michael Broughton, Guillaume Verdon, Trevor McCourt, Antonio J. Martinez, Jae Hyeon Yoo, Sergei V. Isakov, Philip Massey, Ramin Halavati, Murphy Yuezheng Niu, Alexander Zlokapa, Evan Peters, Owen Lockwood, Andrea Skolik, Sofiene Jerbi, Vedran Dunjko, Martin Leib, Michael Streif, David Von Dollen, Hongxiang Chen, Shuxiang Cao, Roeland Wiersema, Hsin-Yuan Huang, Jarrod R. McClean, Ryan Babbush, Sergio Boixo, Dave Bacon, Alan K. Ho, Hartmut Neven, and Masoud Mohseni. “Tensorflow quantum: A software framework for quantum machine learning”. *arXiv eprint* (2020) arXiv:2003.02989.
- [30] Ville Bergholm, Josh Izaac, Maria Schuld, Christian Gogolin, M. Sohaib Alam, Shahnawaz Ahmed, Juan Miguel Arrazola, Carsten Blank, Alain Delgado, Soran Jahangiri, Keri McKiernan, Johannes Jakob Meyer, Zeyue Niu, Antal Száva, and Nathan Killoran. “PennyLane: Automatic differentiation of hybrid quantum-classical computations”. *arXiv eprint* (2018) arXiv:1811.04968.
- [31] “Paddle Quantum”. <https://github.com/PaddlePaddle/Quantum> (2020).
- [32] MindQuantum Developer. “Mindquantum, version 0.5.0”. <https://gitee.com/mindspace/mindquantum> (2021).
- [33] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. “Quantum machine learning”. *Nature* **549**, 195–202 (2017).
- [34] Igor L Markov and Yaoyun Shi. “Simulating quantum computation by contracting tensor networks”. *SIAM Journal on Computing* **38**, 963–981 (2008).
- [35] John Brennan, Momme Allalen, David Brayford, Kenneth Hanley, Luigi Iapichino, Lee J. O’Riordan, Myles Doyle, and Niall Moran. “Tensor network circuit simulation at exascale”. *arXiv eprint* (2021) arXiv:2110.09894.

- [36] Eli A. Meiriom, Haggai Maron, Shie Mannor, and Gal Chechik. “Optimizing tensor network contraction using reinforcement learning”. arXiv eprint (2022) arXiv:2204.09052.
- [37] Román Orús. “A practical introduction to tensor networks: Matrix product states and projected entangled pair states”. *Annals of physics* **349**, 117–158 (2014).
- [38] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G.S.L. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naa-man, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. “Quantum supremacy using a programmable superconducting processor”. *Nature* **574**, 505–510 (2019).
- [39] Yulin Wu, Wan-Su Bao, Sirui Cao, Fusheng Chen, Ming-Cheng Chen, Xiawei Chen, Tung-Hsun Chung, Hui Deng, Yajie Du, Daojin Fan, Ming Gong, Cheng Guo, Chu Guo, Shaojun Guo, Lianchen Han, Linyin Hong, He-Liang Huang, Yong-Heng Huo, Liping Li, Na Li, Shaowei Li, Yuan Li, Futian Liang, Chun Lin, Jin Lin, Haoran Qian, Dan Qiao, Hao Rong, Hong Su, Lihua Sun, Liangyuan Wang, Shiyu Wang, Dachao Wu, Yu Xu, Kai Yan, Weifeng Yang, Yang Yang, Yangsen Ye, Jianghan Yin, Chong Ying, Jiale Yu, Chen Zha, Cha Zhang, Haibin Zhang, Kaili Zhang, Yiming Zhang, Han Zhao, Youwei Zhao, Liang Zhou, Qingling Zhu, Chao-Yang Lu, Cheng-Zhi Peng, Xiaobo Zhu, and Jian-Wei Pan. “Strong quantum computational advantage using a superconducting quantum processor”. *Phys. Rev. Lett.* **127**, 180501 (2021).
- [40] Chu Guo, Yong Liu, Min Xiong, Shichuan Xue, Xiang Fu, Anqi Huang, Xiaogang Qiang, Ping Xu, Junhua Liu, Shenggen Zheng, He-Liang Huang, Mingtang Deng, Dario Poletti, Wan-Su Bao, and Junjie Wu. “General-purpose quantum circuit simulator with projected entangled-pair states and the quantum supremacy frontier”. *Phys. Rev. Lett.* **123**, 190501 (2019).
- [41] Feng Pan and Pan Zhang. “Simulation of quantum circuits using the big-batch tensor network method”. *Phys. Rev. Lett.* **128**, 030501 (2022).
- [42] Yong (Alexander) Liu, Xin (Lucy) Liu, Fang (Nancy) Li, Haohuan Fu, Yuling Yang, Jiawei Song, Pengpeng Zhao, Zhen Wang, Dajia Peng, Huarong Chen, Chu Guo, Heliang Huang, Wenzhao Wu, and Dexun Chen. “Closing the "quantum supremacy" gap”. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. ACM (2021).
- [43] Cupjin Huang, Fang Zhang, Michael Newman, Junjie Cai, Xun Gao, Zhengxiong Tian, Junyin Wu, Haihong Xu, Huanjun Yu, Bo Yuan, Mario Szegedy, Yaoyun Shi, and Jianxin Chen. “Classical simulation of quantum supremacy circuits”. arXiv eprint (2020) arXiv:2005.06787.
- [44] Xin Liu, Chu Guo, Yong Liu, Yuling Yang, Jiawei Song, Jie Gao, Zhen Wang, Wenzhao Wu, Dajia Peng, Pengpeng Zhao, Fang Li, He-Liang Huang, Haohuan Fu, and Dexun Chen. “Redefining the quantum supremacy baseline with a new generation sunway supercomputer”. arXiv eprint (2021) arXiv:2111.01066.
- [45] Feng Pan, Keyang Chen, and Pan Zhang. “Solving the sampling problem of the sycamore quantum supremacy circuits”. arXiv eprint (2021) arXiv:2111.03011.
- [46] Chase Roberts, Ashley Milsted, Martin Ganahl, Adam Zalcman, Bruce Fontaine, Yijian Zou, Jack Hidary, Guifre Vidal, and Stefan Leichenauer. “Tensornetwork: A library for physics and machine learning”. arXiv eprint (2019) arXiv:1905.01330.

- [47] Shi-Xin Zhang, Chang-Yu Hsieh, Shengyu Zhang, and Hong Yao. “Differentiable quantum architecture search”. *Quantum Sci. Technol.* **7**, 045023 (2022).
- [48] Shi-Xin Zhang, Chang-Yu Hsieh, Shengyu Zhang, and Hong Yao. “Neural predictor based quantum architecture search”. *Machine Learning: Science and Technology* **2**, 045027 (2021).
- [49] Shi-Xin Zhang, Zhou-Quan Wan, Chee-Kong Lee, Chang-Yu Hsieh, Shengyu Zhang, and Hong Yao. “Variational Quantum-Neural Hybrid Eigensolver”. *Physical Review Letters* **128**, 120502 (2022).
- [50] Shi-Xin Zhang, Zhou-Quan Wan, Chang-Yu Hsieh, Hong Yao, and Shengyu Zhang. “Variational quantum-neural hybrid error mitigation”. arXiv eprint (2021) arXiv:2112.10380.
- [51] Shuo Liu, Shi-Xin Zhang, Chang-Yu Hsieh, Shengyu Zhang, and Hong Yao. “Probing many-body localization by excited-state vqe”. arXiv eprint (2021) arXiv:2111.13719.
- [52] Shuo Liu, Shi-Xin Zhang, Chang-Yu Hsieh, Shengyu Zhang, and Hong Yao. “Discrete time crystal enabled by stark many-body localization” (2022). arXiv:2208.02866.
- [53] Yu-Qin Chen, Shi-Xin Zhang, Chang-Yu Hsieh, and Shengyu Zhang. “A non-hermitian ground state searching algorithm enhanced by variational toolbox”. arXiv eprint (2022) arXiv:2210.09007.
- [54] Alastair Kay. “Tutorial on the quantikz package”. arXiv eprint (2018) arXiv:1809.03842.
- [55] Matteo Hessel, David Budden, Fabio Viola, Mihaela Rosca, Eren Sezener, and Tom Hennigan. “Optax: composable gradient transformation and optimisation, in jax!” (2020).
- [56] Yaodong Li, Xiao Chen, and Matthew P. A. Fisher. “Quantum zeno effect and the many-body entanglement transition”. *Phys. Rev. B* **98**, 205136 (2018).
- [57] Amos Chan, Rahul M. Nandkishore, Michael Pretko, and Graeme Smith. “Unitary-projective entanglement dynamics”. *Phys. Rev. B* **99**, 224307 (2019).
- [58] Brian Skinner, Jonathan Ruhman, and Adam Nahum. “Measurement-induced phase transitions in the dynamics of entanglement”. *Phys. Rev. X* **9**, 031009 (2019).
- [59] Yaodong Li, Xiao Chen, and Matthew P. A. Fisher. “Measurement-driven entanglement transition in hybrid quantum circuits”. *Phys. Rev. B* **100**, 134306 (2019).
- [60] Shuo Liu, Ming-Rui Li, Shi-Xin Zhang, Shao-Kai Jian, and Hong Yao. “Universal kpz scaling in noisy hybrid quantum circuits”. arXiv eprint (2022) arXiv:2212.03901.
- [61] Johnnie Gray. “quimb: A python package for quantum information and many-body calculations”. *Journal of Open Source Software* **3**, 819 (2018).
- [62] Steven R. White. “Density matrix formulation for quantum renormalization groups”. *Phys. Rev. Lett.* **69**, 2863–2866 (1992).
- [63] Daniel G. a. Smith and Johnnie Gray. “opt_einsum - a python package for optimizing contraction order for einsum-like expressions”. *Journal of Open Source Software* **3**, 753 (2018).
- [64] James Stokes, Josh Izaac, Nathan Killoran, and Giuseppe Carleo. “Quantum natural gradient”. *Quantum* **4**, 269 (2020).
- [65] Xiao Yuan, Suguru Endo, Qi Zhao, Ying Li, and Simon C Benjamin. “Theory of variational quantum simulation”. *Quantum* **3**, 191 (2019).
- [66] Suguru Endo, Jinzhao Sun, Ying Li, Simon C. Benjamin, and Xiao Yuan. “Variational quantum simulation of general processes”. *Phys. Rev. Lett.* **125**, 010501 (2020).
- [67] Sam McArdle, Suguru Endo, Alán Aspuru-Guzik, Simon C. Benjamin, and Xiao Yuan. “Quantum computational chemistry”. *Rev. Mod. Phys.* **92**, 015003 (2020).

- [68] Yudong Cao, Jonathan Romero, Jonathan P Olson, Matthias Degroote, Peter D. Johnson, Mária Kieferová, Ian D Kivlichan, Tim Menke, Borja Peropadre, Nicolas P D Sawaya, Sukin Sim, Libor Veis, and Alán Aspuru-Guzik. “Quantum Chemistry in the Age of Quantum Computing”. *Chemical Reviews* **119**, 10856–10915 (2019).
- [69] Jarrod R. McClean, Kevin J. Sung, Ian D. Kivlichan, Yudong Cao, Chengyu Dai, E. Schuyler Fried, Craig Gidney, Brendan Gimby, Pranav Gokhale, Thomas Häner, Tarini Hardikar, Vojtěch Havlíček, Oscar Higgott, Cupjin Huang, Josh Izaac, Zhang Jiang, Xinle Liu, Sam McArdle, Matthew Neeley, Thomas O’Brien, Bryan O’Gorman, Isil Ozfidan, Maxwell D. Radin, Jhonathan Romero, Nicholas Rubin, Nicolas P. D. Sawaya, Kanav Setia, Sukin Sim, Damian S. Steiger, Mark Steudtner, Qiming Sun, Wei Sun, Daochen Wang, Fang Zhang, and Ryan Babbush. “Openfermion: The electronic structure package for quantum computers”. arXiv eprint (2017) arXiv:1710.07629.
- [70] Bernhard Jobst, Adam Smith, and Frank Pollmann. “Finite-depth scaling of infinite quantum circuits for quantum critical points”. *Physical Review Research* **4**, 033118 (2022).
- [71] Martin Larocca, Nathan Ju, Diego García-Martín, Patrick J. Coles, and M. Cerezo. “Theory of overparametrization in quantum neural networks” (2021). arXiv:2109.11676.
- [72] Navin Khaneja, Timo Reiss, Cindie Kehlet, Thomas Schulte-Herbrüggen, and Steffen J. Glaser. “Optimal control of coupled spin dynamics: design of nmr pulse sequences by gradient ascent algorithms”. *Journal of Magnetic Resonance* **172**, 296–305 (2005).
- [73] Xiaotong Ni, Hui-Hai Zhao, Lei Wang, Feng Wu, and Jianxin Chen. “Integrating quantum processor device and control optimization in a gradient-based framework”. *npj Quantum Inf.* **8**, 106 (2022).
- [74] Guifré Vidal. “Efficient classical simulation of slightly entangled quantum computations”. *Phys. Rev. Lett.* **91**, 147902 (2003).
- [75] Yiqing Zhou, E. Miles Stoudenmire, and Xavier Waintal. “What limits the simulation of quantum computers?”. *Phys. Rev. X* **10**, 041038 (2020).
- [76] J.C. Spall. “Adaptive stochastic approximation by the simultaneous perturbation method”. *IEEE Transactions on Automatic Control* **45**, 1839–1853 (2000).
- [77] Mateusz Ostaszewski, Edward Grant, and Marcello Benedetti. “Structure optimization for parameterized quantum circuits”. *Quantum* **5**, 391 (2021).
- [78] I. A. Luchnikov, A. Ryzhov, S. N. Filippov, and H. Ouerdane. “QGOpt: Riemannian optimization for quantum technologies”. *SciPost Phys.* **10**, 79 (2021).