

Universal quantum circuits for quantum chemistry

Juan Miguel Arrazola, Olivia Di Matteo, Nicolás Quesada, Soran Jahangiri, Alain Delgado, and Nathan Killoran

Xanadu, Toronto, ON, M5G 2C8, Canada

Universal gate sets for quantum computing have been known and studied for decades, yet less is understood about universal gate sets for particle-conserving unitaries, which are operations of interest in quantum chemistry. In this work, we show that controlled single-excitation gates in the form of Givens rotations are universal for particle-conserving unitaries. Single-excitation gates describe an arbitrary $U(2)$ rotation on the two-qubit subspace spanned by the states $|01\rangle, |10\rangle$, while leaving other states unchanged – a transformation that is analogous to a single-qubit rotation on a dual-rail qubit. The proof is constructive, so our result also provides an explicit method for compiling arbitrary particle-conserving unitaries. Additionally, we describe a method for using controlled single-excitation gates to prepare an arbitrary state of a fixed number of particles. We derive analytical gradient formulas for Givens rotations as well as decompositions into single-qubit and CNOT gates. Our results offer a unifying framework for quantum computational chemistry where every algorithm is a unique recipe built from the same universal ingredients: Givens rotations.

1 Introduction

Quantum algorithms for quantum chemistry rely on the ability to prepare states that represent fermionic wavefunctions [1, 2, 3]. These can correspond to ground and excited states of molecular Hamiltonians, which can then be employed to compute properties of the molecule [4, 5, 6, 7, 8, 9]. In most molecules and materials, the number of particles is a conserved quantity and quantum states that do not respect this condition are unphysical. Valid quantum states thus occupy only a subspace of the available Hilbert space. Algorithms that access subspaces of different particle number are therefore not only wasteful, but can potentially lead to incorrect outcomes. This motivates the use of gate sets that preserve subspaces of fixed particle number. We focus on the Jordan-Wigner representation [10], which encodes the subspace of states with k particles in n spin-orbitals into n qubits. This

Juan Miguel Arrazola: juanmiguel@xanadu.ai

space is spanned by the set of all n -qubit states with Hamming weight k , i.e., states with k ones and $n - k$ zeros. To ensure that output states remain valid, quantum circuits for quantum chemistry benefit from employing gates that preserve Hamming weight and therefore particle number.

Universal gate sets capable of synthesizing arbitrary unitary operations have been known for decades [11, 12]. Famously, the set of arbitrary single-qubit rotations and CNOT gates is universal for quantum computation [13]. Yet less is known about universal gate sets for particle-conserving unitaries, which are precisely the operations of interest in quantum chemistry. A notable exception is the work of Ref. [14], which provides a non-constructive proof that Gaussian fermionic operations together with a non-Gaussian gate is universal for particle-conserving transformations. Additionally, there are several proposals for preparing states of fermionic systems [15, 16, 17], some of which are designed to preserve particle number and other symmetries [18, 19, 20]. A universal set of particle-conserving gates, together with a constructive method for compiling arbitrary unitaries, would constitute a flexible and composable framework for designing arbitrary quantum circuits for quantum chemistry.

Currently, a wide variety of quantum circuit architectures have been proposed to prepare states of many-body fermionic systems, particularly in the context of variational quantum algorithms. These include chemically-inspired circuits [21, 22], adaptive circuits [23, 24], hardware-efficient circuits [25, 26], and other specialized methods [27, 28]. This situation is not ideal because quantum algorithm developers are seemingly faced with a choice among different proposals rather than having access to universal building blocks to construct any desired operation. There is therefore a need for a unified conceptual framework for constructing and designing quantum circuits for quantum chemistry. This is particularly important for software implementations that aim to provide users with maximum flexibility and to be adaptable in incorporating future algorithmic innovations.

In this work, *we provide such a framework* by giving a constructive proof that controlled single-excitation gates are universal for particle-conserving unitaries. A single-excitation gate performs an arbitrary $U(2)$ transformation in the subspace $|01\rangle, |10\rangle$

while leaving other basis states unchanged:

$$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & a & c & 0 \\ 0 & b & d & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (1)$$

where $U = \begin{pmatrix} a & c \\ b & d \end{pmatrix}$ is a general 2×2 unitary. Single-excitation gates can be viewed as an extension of *Givens rotations* to unitary two-dimensional transformations. A controlled single-excitation gate, which applies this Givens rotation depending on the state of a third qubit, can be described by the unitary

$$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & a & c & 0 \\ 0 & 0 & 0 & 0 & 0 & b & d & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}. \quad (2)$$

In addition to the universality result, we propose an explicit algorithm using excitation gates to prepare an arbitrary state with a fixed number of particles. We derive analytical gradient formulas for Givens rotations and argue that they are ideal building blocks in variational quantum circuits for quantum chemistry. Our results and framework allow the design of quantum algorithms for quantum chemistry from a universal gate set of operations that respect the symmetries of fermionic systems and have a transparent physical interpretation. Although we focus on the specific case of quantum chemistry, our results apply more generally to the simulation of many-body fermionic systems in the Jordan-Wigner representation. As described in Ref. [29], the framework described in this work has been implemented in the open-source library PennyLane, providing actual building blocks for building arbitrary particle-conserving circuits.

The rest of this manuscript is organized as follows. We introduce the basic concepts of particle-conserving unitaries and Givens rotations in Sec. 2. We then show in Sec. 3 that controlled single-excitation gates are universal for particle-conserving unitaries. In Sec. 4 we describe a universal method for preparing states with a fixed number of particles. We discuss the role of Givens rotations in variational quantum circuits in Sec. 5 and conclude in Sec. 6.

2 Universal gate set

We introduce basic concepts and notation that are relevant before presenting the main results. For simplicity and generality, we employ abstract notions of particles and operations, without establishing an explicit connection to concepts such as electrons, spin-orbitals, or fermionic operators.

2.1 Particle-conserving unitaries

Define the qubit ladder operators $\sigma^\dagger = (X + iY)/2$, $\sigma = (X - iY)/2$, where X, Y are Pauli matrices, and the total number operator

$$N = \sum_i \sigma_i^\dagger \sigma_i. \quad (3)$$

For a computational basis state $|x\rangle$ it holds that $N|x\rangle = w(x)|x\rangle$, where $w(x)$ is the Hamming weight of the bit string x . We define the Hamming weight to be equal to the number of particles and refer to eigenstates of the total number operator as states with a fixed number of particles. A unitary gate U is deemed particle-conserving if

$$[U, N] = 0. \quad (4)$$

A particle-conserving unitary maps states with a fixed number of particles to other states with the same fixed number of particles. Any product of particle-conserving unitaries is also particle-conserving, so any quantum circuit composed of particle-conserving gates is guaranteed to perform a particle-conserving transformation.

The space of all states with k particles on n qubits, denoted as \mathcal{H}_k , is spanned by the set of computational basis states with Hamming weight k . In general, any state of a fixed number of particles can be interpreted as an excitation from a reference state. Unless stated otherwise, we use the state $|11 \dots 100 \dots 0\rangle$ with all particles in the first k qubits as the reference state. This is illustrated in Fig 1. We use the Hamming distance $\sum_{i=1}^n x_i \oplus y_i$ to denote the number of qubits where the computational basis states $|x\rangle$ and $|y\rangle$ differ. For states with an equal number of particles, the Hamming distance is an even number. An example of particle-conserving unitaries are fermionic linear-optical transformations as explored in Ref. [14], which together with a non-Gaussian gate form a universal gate set for particle-conserving unitaries.

Two states $|x\rangle$ and $|y\rangle$ having an equal number of particles are said to differ by an excitation of order ℓ if their Hamming distance is equal to 2ℓ . For example, the states $|1100\rangle$ and $|0101\rangle$ differ by a single excitation (order 1) from the first to the fourth qubit. Similarly, the state $|0011\rangle$ differs from $|1100\rangle$ by a double excitation (order 2).

Any particle-conserving unitary acting on states with k particles on n qubits can be represented as a block-diagonal unitary performing a general $U(d)$ transformation on the subspace \mathcal{H}_k with dimension $d = \dim(\mathcal{H}_k)$. For universality, it is therefore sufficient to consider a set of particle-conserving gates that is universal for the subspace \mathcal{H}_k .

2.2 Givens rotations

As discussed above, any two states with a fixed number of particles differ by an excitation of a given or-

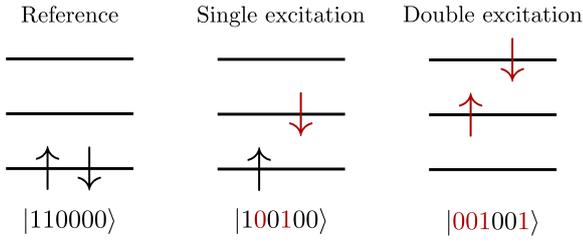


Figure 1: Jordan-Wigner representation of states with a fixed number of particles. Each qubit corresponds to an orbital with a specific spin orientation. The state of the qubit determines whether that spin-orbital is occupied or not. All basis states can be obtained from a reference state, in this case $|110000\rangle$, by a specific excitation. For instance, the state $|100100\rangle$ is obtained by exciting a particle from qubit 2 to 4, while the state $|001001\rangle$ is obtained by exciting both particles to qubits 3 and 4.

der. It is therefore convenient to work with a set of quantum gates that create superpositions between the original and the excited state. In the simplest non-trivial case of a single particle and two qubits, these correspond to gates that perform arbitrary $U(2)$ rotations between the states $|10\rangle, |01\rangle$ while leaving other basis states unchanged. For example, restricting to the case where the gate has only real parameters, a two-qubit particle-conserving unitary can be written as

$$G(\theta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta) & -\sin(\theta) & 0 \\ 0 & \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (5)$$

where we use the ordering $|00\rangle, |01\rangle, |10\rangle, |11\rangle$ of two-qubit computational basis states.

This is an example of a *Givens rotation*: a rotation in a two-dimensional subspace of a larger space. In this case, the Givens rotation acts as a two-qubit *single-excitation gate*, coupling states that differ by a single excitation. More generally, we can extend the concept of Givens rotations to $U(2)$ transformations in two-dimensional subspaces, where a general single-excitation gate can be written as

$$G = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & a & c & 0 \\ 0 & b & d & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (6)$$

where $|a|^2 + |c|^2 = |b|^2 + |d|^2 = 1$ and $ab^* + cd^* = 0$ to ensure unitarity. We can also consider four-qubit *double-excitation gates* $G^{(2)}$ which perform a general $U(2)$ rotation on the subspace spanned by the states $|0011\rangle, |1100\rangle$

$$G^{(2)} |0011\rangle = a |0011\rangle + b |1100\rangle, \quad (7)$$

$$G^{(2)} |1100\rangle = d |1100\rangle + c |0011\rangle, \quad (8)$$

while leaving all remaining four-qubit states unchanged. Double-excitation gates can also perform

rotations in two-dimensional subspaces defined by pairs of four-qubit states with Hamming distance four, namely $|1010\rangle, |0101\rangle$ and $|1001\rangle, |0110\rangle$.

We can generalize to excitation gates of order ℓ . These are unitary Givens rotations acting on the space of 2ℓ qubits that couple the states $|\mathbf{1}_\ell \mathbf{0}_\ell\rangle := |1\rangle^{\otimes \ell} |0\rangle^{\otimes \ell}$ and $|\mathbf{0}_\ell \mathbf{1}_\ell\rangle := |0\rangle^{\otimes \ell} |1\rangle^{\otimes \ell}$ as

$$G^{(\ell)} |\mathbf{0}_\ell \mathbf{1}_\ell\rangle = a |\mathbf{0}_\ell \mathbf{1}_\ell\rangle + b |\mathbf{1}_\ell \mathbf{0}_\ell\rangle, \quad (9)$$

$$G^{(\ell)} |\mathbf{1}_\ell \mathbf{0}_\ell\rangle = d |\mathbf{1}_\ell \mathbf{0}_\ell\rangle + c |\mathbf{0}_\ell \mathbf{1}_\ell\rangle, \quad (10)$$

while acting as the identity on all other states. Similar Givens rotations can be defined for permutations of the states $|\mathbf{1}_\ell \mathbf{0}_\ell\rangle, |\mathbf{0}_\ell \mathbf{1}_\ell\rangle$, i.e., excitation gates of order ℓ also include rotations on all pairs of states of 2ℓ qubits with Hamming distance 2ℓ . By construction, these excitation gates are particle-conserving since they only couple states with an equal number of particles.

As explained in more detail in Sec. 3, on their own, excitation gates as Givens rotations are not universal for particle-conserving unitaries. We thus consider *controlled* excitation gates, which apply an excitation gate depending on the state of a control qubit. In particular, we focus on the three-qubit controlled single-excitation gate

$$CG = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & a & c & 0 \\ 0 & 0 & 0 & 0 & 0 & b & d & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}. \quad (11)$$

A particular example of a controlled single-excitation gate is the controlled SWAP, or Fredkin gate

$$F = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}. \quad (12)$$

Controlled gates usually refer to the case where a gate is applied only if the control qubit is in state $|1\rangle$. Throughout this work, we more generally use the term controlled gate to include also the case where gates are applied only if the control qubit is in state $|0\rangle$. All such controlled gates are particle-conserving. Controlled single-excitation gates and double-excitation gates are illustrated in Fig. 2. It is also useful to note that Givens rotations can also be used to enforce spin conservation by coupling only qubits corresponding to spin-orbitals with the same spin projection. This allows a straightforward method for ensuring that spin-conservation is also respected in a quantum circuit.

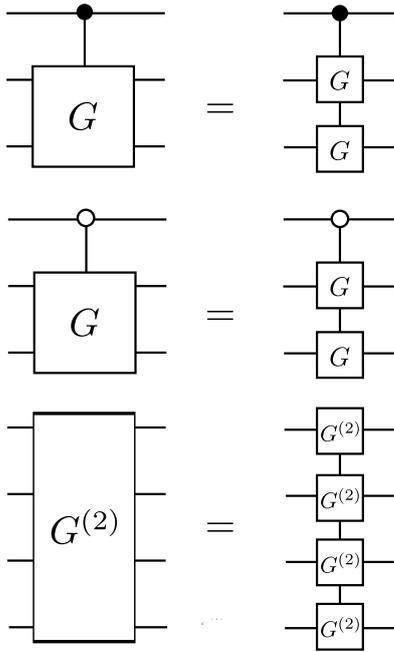


Figure 2: The three-qubit controlled single-excitation gate, which is universal for particle-conserving unitaries, and the four qubit double-excitation gate. The controlled single-excitation gate on the top denotes a control on state $|1\rangle$ while the one below shows a control on state $|0\rangle$. The figures on the right are alternative representations of the excitation gates, which are useful for drawing circuits when the qubits are not adjacent. The bottom circuit is a double-excitation gate.

It is helpful to contrast Givens rotations to fermionic excitations, which are operators of the form $\exp[-i\theta(a_i^\dagger a_j - a_j^\dagger a_i)]$ for a single excitation, where a, a^\dagger are fermionic ladder operators. Higher-order excitations are constructed with higher-order monomials of the ladder operators. Fermionic excitations are expressed as qubit operators through the Jordan-Wigner mapping $a_i^\dagger \rightarrow \sigma_i^+ \prod_{j<i} Z_j$, where $\sigma^\pm = (X \pm iY)$. Fermionic excitations also couple subspaces of equal particle number, but they are more difficult to analyze, implement, and simulate. The universality proof of the next section demonstrates that it is possible to instead employ Givens rotations directly since, as we will show later in Sec. 4, they can be used to prepare any desired state.

3 Proof of Universality

In this section, we show that controlled single-excitation gates are universal for particle-conserving unitaries. We use standard textbook methods similar to those used for proving the universality of single-qubit and CNOT gates [13]. The proof follows these main steps:

1. For particle-conserving unitaries, we show that the relevant $U(2)$ transformations are excitation

gates controlled on multiple qubits. Using the established result that $U(d)$ transformations can be decomposed into products of $U(2)$ transformations, it follows that excitation gates controlled on multiple qubits are universal for particle-conserving unitaries.

2. We show that any excitation gate controlled on multiple qubits can be decomposed in terms of multiply-controlled single-excitation gates.
3. Finally, we show that single excitation gates controlled on multiple qubits can be decomposed into three-qubit controlled single-excitation gates, which are therefore also universal.

The main insight of this construction is that single-excitation gates are analogous to single-qubit gates. Indeed, the subspace $|10\rangle, |01\rangle$ can be interpreted as a dual-rail encoding of a single qubit. This allows standard methods to carry over to the particle-conserving case, with few modifications. Fredkin gates, which are a specific type of control single-excitation gate, play a special role: they are used to extend controlled single-excitation gates to controlled gates over multiple qubits. Fredkin gates have been shown to be universal for reversible computations in dual-rail encodings [30].

3.1 Excitation gates with multiple controls

It is well-established that any $U(d)$ transformation can be decomposed into a product of $U(2)$ transformations acting on arbitrary two-dimensional subspaces [31, 13, 32, 33]. As discussed in the previous section, any state of a fixed number of particles can be obtained by applying an excitation to a reference state. This result holds more generally: any two states of a fixed number of particles differ by a specific excitation.

Consider two k -particle states on n qubits $|x\rangle, |y\rangle$, with Hamming distance 2ℓ . Without loss of generality, since this can be guaranteed by relabelling, suppose the first $k - \ell$ qubits are set to 1 for both states and the last $n - k - \ell$ qubits are set to 0. The remaining 2ℓ qubits are in different states, meaning that they can be mapped to each other by exciting the particles from the ℓ occupied qubits to the ℓ unoccupied ones.

For example, the states $|111000\rangle$ and $|110010\rangle$, which have Hamming distance 2, differ by a single excitation. Similarly, the states $|011010\rangle$ and $|010101\rangle$ differ by a double excitation. This connection between states and excitations is illustrated in Fig. 3.

A $U(2)$ rotation in the subspace spanned by the k -particle states $|x\rangle, |y\rangle$ is therefore equivalent to a

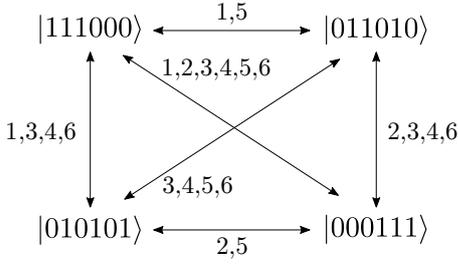


Figure 3: Any pair of states with a fixed number of particles can be related by an excitation. The states $|111000\rangle$ and $|011010\rangle$ differ by a single excitation between qubits 1 and 5. The states $|011010\rangle$ and $|000111\rangle$ differ by a double excitation between qubits 2,3,4,6, while $|000111\rangle$ and $|111000\rangle$ are connected by a triple excitation acting on all qubits.

unitary performing the transformation

$$U|x\rangle = a|x\rangle + b|y\rangle, \quad (13)$$

$$U|y\rangle = d|y\rangle + c|x\rangle, \quad (14)$$

while leaving every other basis state unchanged. For the specific case of states with $n = 2\ell$ qubits, this is accomplished by a unitary Givens rotation $G^{(\ell)}$, as in Eqs. (9) and (10). When applied to states with $n > 2\ell$ qubits, the gate $G^{(\ell)}$ acts non-trivially on any states where the 2ℓ qubits in question are set to $|x\rangle$ or $|y\rangle$, regardless of the state of the remaining qubits. For instance, if $|z\rangle$ is a basis state of $m = n - 2\ell$ qubits, it holds that

$$G^{(\ell)}|z\rangle|x\rangle = a|z\rangle|x\rangle + b|z\rangle|y\rangle, \quad (15)$$

for all z . To address this issue, we can simply apply $G^{(\ell)}$ controlled on the state of the remaining $n - 2\ell$ qubits. We use the notation $C^{(m)}G^{(\ell)}$ to denote a $G^{(\ell)}$ gate controlled on the state of m qubits. Controlling on the remaining qubits being in state $|z^*\rangle$ and defining $|x'\rangle = |z^*\rangle|x\rangle$ and $|y'\rangle = |z^*\rangle|y\rangle$, we then have that

$$C^{(m)}G^{(\ell)}|x'\rangle = a|x'\rangle + b|y'\rangle, \quad (16)$$

$$C^{(m)}G^{(\ell)}|y'\rangle = d|y'\rangle + c|x'\rangle, \quad (17)$$

while leaving all other basis states unchanged. This is the desired two-dimensional transformation.

Consider the states $|100011\rangle$ and $|010011\rangle$. They differ by an excitation from the first to the second qubit, and coincide on the remaining four qubits. A non-controlled single-excitation gate acting on the first two qubits would also perform a transformation on other subspaces, for example the one spanned by $|101101\rangle$ and $|011100\rangle$. However, controlling on the last four qubits being in state $|0011\rangle$ ensures that the gate $C^{(4)}G$ acts non-trivially only on the target two-dimensional subspace. The role of multiple controls is shown in Fig. 4.

Overall, we conclude that any two-level $U(2)$ gate on the subspace of k -particle states on n qubits can be

$$C^{(4)}G|10001\rangle = a|10001\rangle + b|01001\rangle$$

$$C^{(4)}G|01001\rangle = |01001\rangle$$

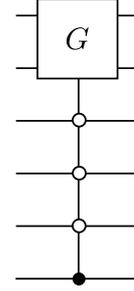


Figure 4: Circuit diagram of a multi-controlled single-excitation gate. A single-excitation gate G can act non-trivially on many states. To ensure that the desired rotation happens only between two target states, we can apply a multi-controlled gate $C^{(4)}G$ that acts as the identity on non-target states. In this example, the goal is to perform a rotation in the subspace of states $|10001\rangle, |01001\rangle$. This can be achieved by applying a single-excitation gate to the first two qubits, controlled on the state of the last four being $|0001\rangle$. This guarantees that other states, such as $|01001\rangle$, are left unchanged.

implemented in terms of multi-controlled excitation gates $C^{(m)}G^{(\ell)}$, where $n = 2\ell + m$. Following standard results, this implies that multi-controlled excitation gates are universal for particle-conserving operations.

3.2 Single-excitation gates with multiple controls

We now show that multi-controlled excitation gates can be decomposed in terms of multi-controlled single-excitation gates $C^{(m)}G$. The construction follows a similar approach to Ref. [13]. Suppose we wish to decompose a $U(2)$ Givens rotation on the subspace spanned by $|x\rangle, |y\rangle$, where the states have Hamming distance 2ℓ . The goal is to employ single-excitation gates to perform a permutation of all basis states such that the permuted versions of $|x\rangle$ and $|y\rangle$ differ by a single excitation. This can be achieved following similar principles to the construction of Gray codes. A controlled single-excitation gate can then be applied, followed by a reversal of the permutation.

For example, the states $|101001\rangle$ and $|010110\rangle$, which differ by a triple excitation, can be linked through the sequence $|101001\rangle \rightarrow |011001\rangle \rightarrow |010101\rangle \rightarrow |010110\rangle$, where each new state differs from the previous one by a single excitation. Each of the states in this sequence can be obtained from the previous one by applying a SWAP gate

$$\text{SWAP} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (18)$$

$$|g_1\rangle = |x\rangle = |101001\rangle$$

$$C^{(4)}\text{SWAP}|g_1\rangle = |011001\rangle = |g_2\rangle$$

$$C^{(4)}\text{SWAP}|g_2\rangle = |010101\rangle = |g_3\rangle$$

$$C^{(4)}G|g_3\rangle = a|g_3\rangle + b|010110\rangle = a|g_3\rangle + b|y\rangle$$

$$C^{(4)}\text{SWAP}(a|g_3\rangle + b|y\rangle) = a|g_2\rangle + b|y\rangle$$

$$C^{(4)}\text{SWAP}(a|g_2\rangle + b|y\rangle) = a|x\rangle + b|y\rangle$$

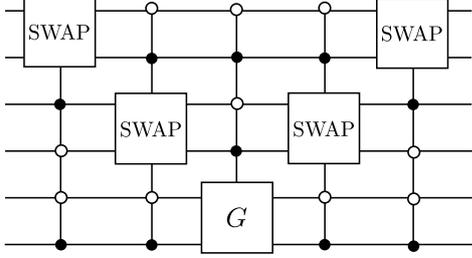


Figure 5: Method for connecting any pair of states by a sequence of particle-conserving multi-controlled SWAP gates. The initial state $|x\rangle = |101001\rangle$ differs from the target state $|y\rangle = |010110\rangle$ by a triple excitation, which can be decomposed in terms of single excitations using a sequence of intermediary states. A circuit implementing a triple-excitation rotation can then be decomposed in terms of (i) multi-controlled SWAP gates performing a permutation of states, (ii) a multi-controlled single-excitation gate, and (iii) a reversal of the permutation. The end result is to create an arbitrary superposition $a|x\rangle + b|y\rangle$ of the target states.

controlled on the state of all remaining qubits, i.e., by applying a $C^{(m)}\text{SWAP}$ gate. The SWAP gate is a special case of a single-excitation gate. As before, the control is required to ensure that the resulting permutation happens non-trivially only on the two-dimensional target subspace. This procedure is illustrated in Fig. 5.

We now describe the method in more detail. Without loss of generality, suppose that the states $|x\rangle$ and $|y\rangle$ differ on the first 2ℓ qubits. The first step is to outline an ordered sequence of computational basis states $|g_1\rangle, |g_2\rangle, \dots, |g_{\ell+1}\rangle$ such that all $|g_i\rangle, |g_{i+1}\rangle$ differ by a single excitation, and where $|x\rangle = |g_1\rangle$ and $|y\rangle = |g_{\ell+1}\rangle$. To build the circuit implementing the decomposition, we perform the following steps.

1. Apply a SWAP gate to the qubits where $|x\rangle = |g_1\rangle$ and $|g_2\rangle$ differ, controlled on all other qubits. This has the effect of swapping $|g_1\rangle, |g_2\rangle$, while leaving all other states unchanged.
2. Follow the same procedure to swap $|g_2\rangle, |g_3\rangle$, then $|g_3\rangle, |g_4\rangle$, and all other states until the final swap between $|g_{\ell-1}\rangle, |g_\ell\rangle$. This sequence of operations has the effect of mapping $|x\rangle \rightarrow |g_\ell\rangle$ while $|g_{\ell+1}\rangle = |y\rangle$ is left unchanged.
3. Since $|g_\ell\rangle, |g_{\ell+1}\rangle$ differ by a single excitation, we perform a multi-controlled single-excitation gate

that acts only on the subspace spanned by these states.

4. The circuit is completed by reverting all the swaps such that the resulting transformation is a rotation in the subspace spanned by $|x\rangle$ and $|y\rangle$, as desired.

3.3 Controlled single-excitation gates are universal

Given a gate controlled on a single qubit, there exist well-established methods to extend the control to additional qubits [13]. Suppose we want to control the operation on the state of m qubits $|z_1 z_2 \dots z_m\rangle$. The strategy relies on employing $m - 1$ auxiliary qubits and Toffoli gates (controlled-CNOT gates), as shown in Fig. 6.

Toffoli gates are not particle-conserving, but they can be decomposed in terms of particle-conserving Fredkin gates by replacing the auxiliary qubits with dual-rail qubits $|\tilde{0}\rangle := |01\rangle, |\tilde{1}\rangle := |10\rangle$. In this case a Toffoli gate is equivalent to a controlled-controlled-SWAP gate, since swapping $|01\rangle$ and $|10\rangle$ applies a NOT gate to the dual-rail qubit. As shown in Ref. [30], the controlled-controlled-SWAP gate can then be decomposed into three Fredkin (controlled-SWAP) gates with the help of an auxiliary qubit, as shown in Fig. 7.

Overall, we have shown that single-excitation gates controlled on multiple qubits, which were previously shown to be universal, can be decomposed into controlled single-excitation gates, which are therefore also universal. This concludes the proof that controlled single-excitation gates are universal for particle-conserving unitaries.

4 State preparation

Universal gate sets for particle-conserving unitaries can also be used to prepare arbitrary states of a fixed number of particles. Here we discuss how controlled single-excitation gates can be used for this purpose. We follow the strategy of Ref. [34].

Consider a system of k particles on n qubits, spanning a space of dimension $d = \binom{n}{k}$. Any such state can be written as $|\psi\rangle = \sum_x c_x |x\rangle$, where the sum is over all n -bit strings x of Hamming weight k . As shown before, an arbitrary $U(2)$ rotation in the subspace of any pair of states $|x\rangle, |y\rangle$ can be performed by a suitable decomposition into controlled single-excitation gates.

Consider a lexicographical labelling of all bit strings of Hamming weight k as $|x_1\rangle, |x_2\rangle, \dots, |x_d\rangle$. For instance, in the case of $n = 3$ and $k = 2$ we have $|x_1\rangle = |011\rangle, |x_2\rangle = |101\rangle$, and $|x_3\rangle = |110\rangle$. An

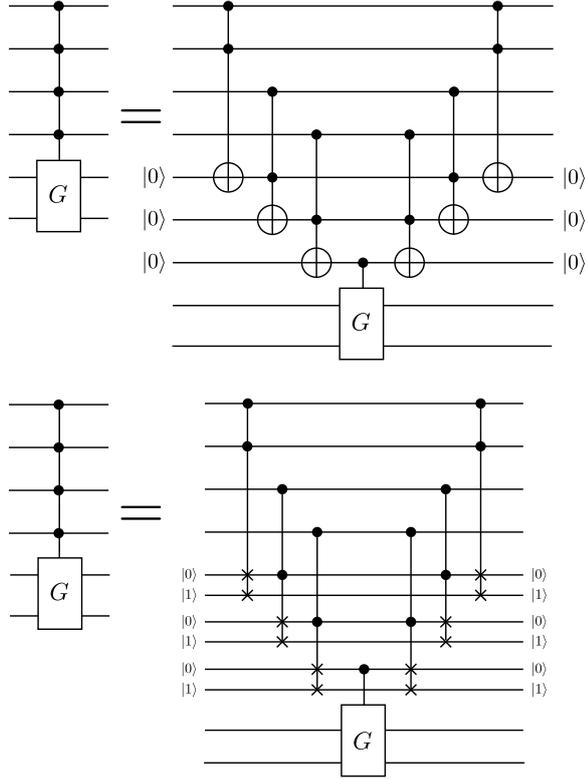


Figure 6: A multi-controlled single-excitation gate can be decomposed in terms of a controlled single-excitation gate using a cascade of Toffoli gates. The Toffoli gates can be decomposed in terms of the particle-conserving controlled-controlled-SWAP gate using additional dual-rail qubits.

arbitrary state can then be written as

$$|\psi\rangle = \sum_{i=1}^d c_i |x_i\rangle. \quad (19)$$

We describe a method to prepare any such state starting from the reference state $|x_1\rangle$.

First, apply the multi-controlled excitation operation in the subspace $|x_1\rangle, |x_2\rangle$ that performs the mapping

$$|x_1\rangle \rightarrow \alpha_1 |x_1\rangle + c_2 |x_2\rangle, \quad (20)$$

where $\alpha_1 = \sqrt{1 - |c_2|^2}$. Then, apply the multi-controlled excitation operation in the subspace $|x_1\rangle, |x_3\rangle$. This performs the mapping

$$\alpha_1 |x_1\rangle + c_2 |x_2\rangle \rightarrow \alpha_1 \alpha_2 |x_1\rangle + c_2 |x_2\rangle + \alpha_1 c'_3 |x_3\rangle, \quad (21)$$

where we set $c'_3 = c_3/\alpha_1$ and $\alpha_2 = \sqrt{1 - |c'_3|^2}$. This ensures that the coefficient in front of $|x_3\rangle$ is precisely the desired one, c_3 . This process can be repeated for each of the remaining states $|x_4\rangle, \dots, |x_d\rangle$. The result is to prepare the state

$$|\psi\rangle = \alpha |x_1\rangle + \sum_{i=2}^d c_i |x_i\rangle, \quad (22)$$

where $\alpha = \prod_{i=1}^{d-1} \alpha_i$. This state is normalized, which from Eq. (19) implies that $|\alpha|^2 = |c_1|^2$. To ensure

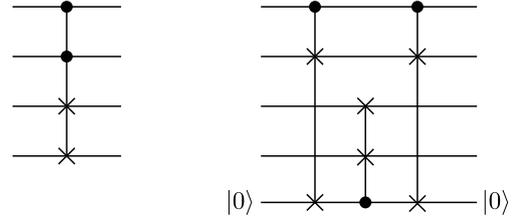


Figure 7: A controlled-controlled-SWAP gate can be decomposed in terms of three Fredkin gates using an auxiliary qubit [30].

that in fact $\alpha = c_1 = |c_1|e^{i\theta}$, it suffices to choose $\alpha_1, \alpha_2, \dots, \alpha_{d-2}$ to be positive real numbers, as we have done, and set $\alpha_{d-1} = |\alpha_{d-1}|e^{i\theta}$ to prepare the desired state. Note that for a unitary $U = \begin{pmatrix} a & c \\ b & d \end{pmatrix}$, a can be guaranteed to be real for any c by choosing

$$d = -\sqrt{1 - |c|^2}$$

$$b^* = \frac{c\sqrt{1 - |c|^2}}{a}.$$

In practice, it is possible to simplify this general strategy when applied to particular cases. Since the gates act only on a specific superposition of basis states, controls only need to be applied on qubits where the states in the superposition differ. This is useful if the target state has support only on a specific subspace. For example, the first excitation gate performing the mapping $|x_1\rangle \rightarrow \alpha_1 |x_1\rangle + c_2 |x_2\rangle$ does not need to be controlled. Furthermore, excitation gates can be chosen to act on different reference states in order to create new superpositions.

Consider the six-qubit state

$$c_1 |110000\rangle + c_2 |001100\rangle + c_3 |000011\rangle + c_4 |100100\rangle, \quad (23)$$

which corresponds to a superposition of the four basis states that contribute most significantly to the ground-state energy of the H_3^+ molecule in a minimal basis set. The state can be prepared as follows.

Starting from $|110000\rangle$, apply a double-excitation gate to the first four qubits to prepare the state $a_1 |110000\rangle + c_2 |001100\rangle$. This does not need to be controlled on any qubit. Then apply a double-excitation gate to qubits 3,4,5,6 to prepare the state $a_1 |110000\rangle + c_2 |001100\rangle + c_3 |000011\rangle$, where we're using $|001100\rangle$ as the reference. This again does not need to be controlled. To obtain the desired state, apply a single-excitation gate to qubits 2 and 4, controlled on the first qubit being in state $|1\rangle$, which prevents mixing with the state $|001100\rangle$. This construction is shown in Fig. 8.

5 Variational quantum circuits

We discuss implications of our results for variational quantum circuits. Our universality result suggests the

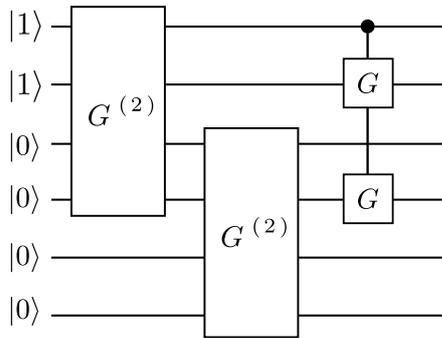


Figure 8: A quantum circuit for preparing the state $c_1 |110000\rangle + c_2 |001100\rangle + c_3 |000011\rangle + c_4 |101000\rangle$ for arbitrary values of the coefficients c_1, c_2, c_3, c_4 . Double-excitation gates are applied to the first four qubits and then to the last four qubits. Finally, a controlled single-excitation gate is applied to qubits 2 and 4, controlled on the state of the first qubit.

use of controlled single-excitation gates as building blocks for variational quantum circuits. For example, the state preparation algorithm described above can be employed as a template where the rotation angles for each gate are free parameters of the model. In this context, multiple controls are not necessary; instead, by employing uncontrolled excitation gates it is possible to reach a larger subspace of states using fewer gates, but generally this makes it more challenging to prepare specific target states. In the example of Fig. 8, dropping the control on the last single excitation gate leads to a state with a non-zero coefficient on the additional basis state $|011000\rangle$. Examples of variational circuits designed using Givens rotations as building blocks are shown in Fig. 9. Below we describe two specific strategies for building quantum circuits for quantum chemistry applications, and derive analytical gradient formulas for Givens rotations.

5.1 All singles and doubles

In the context of quantum computing, the unitary coupled-cluster singles and doubles (UCCSD) ansatz is often expressed in terms of fermionic operators, which are then mapped to complicated qubit gates. In similar spirit to the qubit-coupled-cluster approach of Ref. [27], we can instead consider a circuit where single and double excitations are respectively implemented using Givens rotations G and $G^{(2)}$. A quantum circuit can then be defined consisting of all possible single and double excitation gates that act non-trivially on the reference state without flipping the spin of the excited particles. The resulting circuit is analogous to a Trotterized implementation of UCCSD to first level, but where all gates are Givens rotations. This is illustrated in Fig. 10.

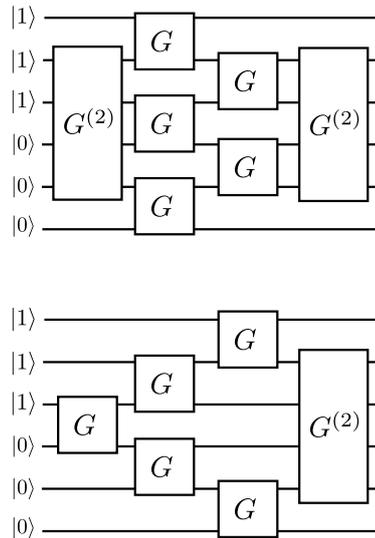


Figure 9: Example circuit architectures constructed from Givens rotations as fundamental building blocks. Since these operations are particle-conserving, it is possible to compose them arbitrarily to create various types of particle-conserving circuits. Excitation gates without controls are used to access larger subspaces with fewer gates, whose parameters may then be optimized for specific purposes.

5.2 Adaptive circuits

Adaptive strategies such as those presented in Refs. [23, 24] can be implemented by selecting Givens rotations instead of fermionic excitations in the constructions. The main idea is that instead of designing circuits that work well for all molecules, we can instead build specific circuits that are custom-built for each molecule. Hence, what is general is the method for building custom circuits, not the circuits themselves.

A simple yet powerful strategy is to build a circuit consisting of all double and single excitation gates, randomly initialize all parameters, and compute the gradient for each gate. The final circuit is constructed by keeping only those gates such that the norm of their gradient exceeds a fixed threshold. This is shown in Fig. 10.

5.3 Analytical gradients

We derive analytic gradient formulas for Givens rotations. If \tilde{H} is the generator of a unitary $\tilde{U}(\theta) = e^{i\tilde{H}\theta}$, then the generator of the unitary $U(\theta) = \mathbb{1} \oplus \tilde{U}(\theta)$ is $H = \mathbf{0} \oplus \tilde{H}$, where $\mathbf{0}$ denotes the zero operator. As shown in [35], generators of this form can be decomposed as

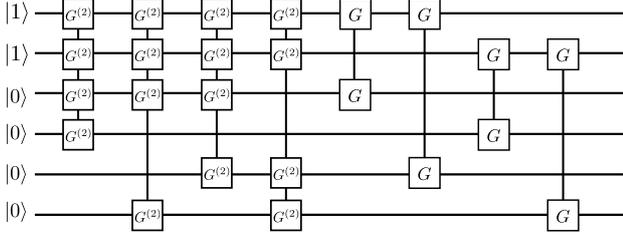
$$H = \frac{1}{2}(H_+ + H_-), \quad (24)$$

$$H_{\pm} = (\pm \mathbb{1}) \oplus \tilde{H}. \quad (25)$$

We can then write

$$U(\theta) = e^{i\theta G_+/2} e^{i\theta G_-/2}, \quad (26)$$

All singles and doubles



Adaptive

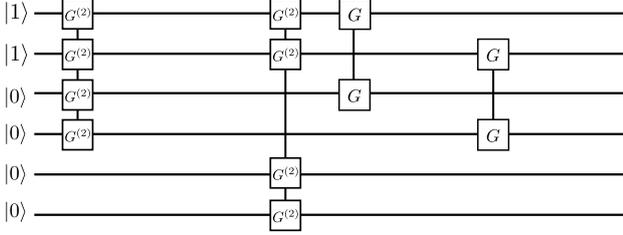


Figure 10: Strategies for building quantum circuits for quantum chemistry. The first strategy consists of selecting all single and double excitation gates that do not flip the spin of the particles. In an adaptive strategy, after initializing gate parameters, we compute the gradient for each gate and keep only those that are above a given threshold.

and define the gates

$$U_{\pm}(\theta) = e^{i\theta G_{\pm}}. \quad (27)$$

The operators H_{\pm} satisfy

$$H_{\pm}^2 = \mathbb{1}, \quad (28)$$

$$[H_+, H_-] = 0. \quad (29)$$

As shown in [36], any unitary $U(\theta)$ with generator that is self-inverse satisfies the parameter-shift rule

$$\frac{\partial C(\theta)}{\partial \theta} = \frac{C(\theta + s) - C(\theta - s)}{2 \sin(s)}, \quad (30)$$

for any cost function that can be written as $C(\theta) = \langle \psi | U^{\dagger}(\theta) K U(\theta) | \psi \rangle$, where K is an observable.

This parameter-shift rule applies to the gates $U_{\pm}(\theta)$, which means that derivatives of $U(\theta)$ can be obtained by writing $U(\theta) = U_+(\theta/2)U_-(\theta/2)$ and computing derivatives of the gates $U_{\pm}(\theta)$. This technique can be employed for any Givens rotation whose generator \tilde{H} is self-inverse. For example, in the case of the Givens rotation of Eq. (5), we can write $G(\theta) = G_+(\theta/2)G_-(\theta/2)$ where the unitaries

$$G_{\pm}(\theta) = \begin{pmatrix} e^{\pm i\theta} & 0 & 0 & 0 \\ 0 & \cos \theta & -\sin \theta & 0 \\ 0 & \sin \theta & \cos \theta & 0 \\ 0 & 0 & 0 & e^{\pm i\theta} \end{pmatrix}, \quad (31)$$

satisfy the parameter-shift rule.

6 Conclusion

We have shown that controlled single-excitation gates are universal for particle-conserving unitaries. These three-qubit gates are Givens rotations performing a transformation in a two-dimensional subspace of states $|01\rangle, |10\rangle$, controlled on the state of a third qubit. The states $|01\rangle, |10\rangle$ can be interpreted as a dual-rail encoding of a single qubit, thus making controlled single-excitation gates analogous to controlled single-qubit gates. This provides a unified approach for combining techniques from quantum computing into a form that is directly applicable to quantum chemistry: instead of single-qubit gates and controlled two-qubit gates, we can focus on single-excitation Givens rotations and controlled Givens rotations. This connection could potentially open avenues for more efficient compilation strategies that employ Givens rotations as the universal gate set while employing standard techniques applicable for general circuits.

The proof of universality relies on the ability to control excitation gates on multiple qubits. This leads to decompositions employing auxiliary qubits and controlled single-excitation gates. These constructions are helpful proof techniques, but may not be the optimal approaches to compiling circuits. Instead, quantum circuits can be designed using Givens rotations directly. The fact that controlled Givens rotations are universal and can be freely composed, while satisfying the main symmetries of fermionic systems, allows for a versatile approach to quantum algorithm design. For example, Givens rotations can also be used to directly enforce spin conservation by coupling only spin-orbitals with equal spin projection. More generally, they can be composed to perform arbitrary transformations on any desired subspace of the fixed-particle-number manifold. For all such circuits, we can obtain analytical gradient rules that can be used to optimize circuits when desired. Overall, our results indicate that Givens rotations are better abstractions than fermionic excitation gates.

For quantum chemistry applications, it is likely that custom algorithms will be needed to tackle specific problems and molecules. Instead of preparing a menu of quantum circuits and algorithms for each instance, we aim to provide scientists with a set of universal ingredients that can be employed to craft tailored solutions. Our results serve as a unifying framework for quantum computational chemistry where every algorithm is a unique recipe built from the same universal ingredients: Givens rotations.

A Decompositions of excitation operators

While controlled single excitation operators comprise a universal gate set, hardware constraints often re-

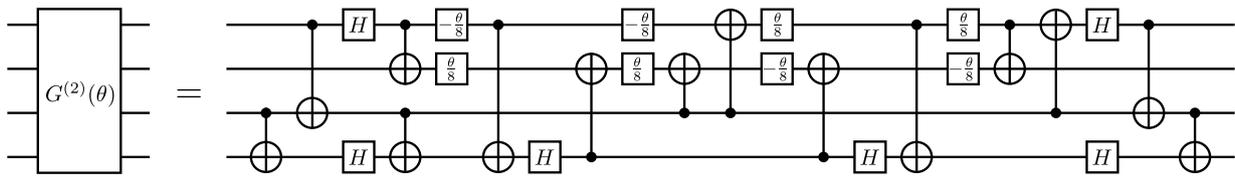


Figure 11: Decomposition of a double excitation gate into single-qubit rotations and CNOTs. All gates denoted by $\pm\frac{\theta}{8}$ are Pauli Y rotations. Adapted from [20].

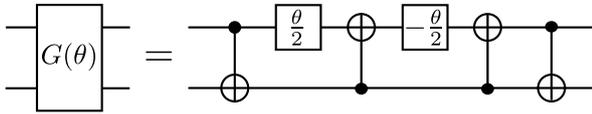


Figure 12: Decomposition of a single excitation gate into single-qubit rotations and CNOTs. The gates denoted by $\pm\frac{\theta}{2}$ are Pauli Y rotations. The middle four gates constitute a controlled Y rotation.

quire a decomposition of these operators over the gate set of single-qubit rotations and CNOTs. Fig. 12 presents such a decomposition for a single-excitation gate. It is straightforward to extend this to a controlled version by applying a control to each individual gate. This would produce Toffolis and controlled-Y rotations, both of which can be further decomposed over the gate set. Fig. 11 presents a decomposition for the four-qubit double-excitation gate. This decomposition was adapted from that of Ref. [20].

References

- [1] Sam McArdle, Suguru Endo, Alan Aspuru-Guzik, Simon C Benjamin, and Xiao Yuan. “Quantum computational chemistry”. *Reviews of Modern Physics* **92**, 015003 (2020).
- [2] Yudong Cao, Jonathan Romero, Jonathan P Olson, Matthias Degroote, Peter D Johnson, Mária Kieferová, Ian D Kivlichan, Tim Menke, Borja Peropadre, Nicolas PD Sawaya, et al. “Quantum chemistry in the age of quantum computing”. *Chemical Reviews* **119**, 10856–10915 (2019).
- [3] Markus Reiher, Nathan Wiebe, Krysta M Svore, Dave Wecker, and Matthias Troyer. “Elucidating reaction mechanisms on quantum computers”. *Proceedings of the National Academy of Sciences* **114**, 7555–7560 (2017).
- [4] 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**, 1–7 (2014).
- [5] Peter JJ O’Malley, Ryan Babbush, Ian D Kivlichan, Jonathan Romero, Jarrod R McClean, Rami Barends, Julian Kelly, Pedram Roushan, Andrew Tranter, Nan Ding, et al. “Scalable quantum simulation of molecular energies”. *Physical Review X* **6**, 031007 (2016).
- [6] Jarrod R McClean, Mollie E Kimchi-Schwartz, Jonathan Carter, and Wibe A De Jong. “Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states”. *Physical Review A* **95**, 042308 (2017).
- [7] Ilya G Ryabinkin, Scott N Genin, and Artur F Izmaylov. “Constrained variational quantum eigensolver: Quantum computer search engine in the fock space”. *Journal of Chemical Theory and Computation* **15**, 249–255 (2018).
- [8] Kosuke Mitarai, Yuya O Nakagawa, and Wataru Mizukami. “Theory of analytical energy derivatives for the variational quantum eigensolver”. *Physical Review Research* **2**, 013129 (2020).
- [9] Yohei Ibe, Yuya O. Nakagawa, Nathan Earnest, Takahiro Yamamoto, Kosuke Mitarai, Qi Gao, and Takao Kobayashi. “Calculating transition amplitudes by variational quantum deflation”. *Phys. Rev. Research* **4**, 013173 (2022).
- [10] P Jordan and E Wigner. “Über das paulische äquivalenzverbot”. *Zeitschrift für Physik* **47**, 631–651 (1928).
- [11] Seth Lloyd. “Almost any quantum logic gate is universal”. *Physical Review Letters* **75**, 346 (1995).
- [12] David P DiVincenzo. “Two-bit gates are universal for quantum computation”. *Physical Review A* **51**, 1015 (1995).
- [13] Michael A Nielsen, Isaac L Chuang, Isaac L Chuang, et al. “Quantum computation and quantum information”. *Cambridge University Press*. (2000).
- [14] Michał Oszmaniec and Zoltán Zimborás. “Universal extensions of restricted classes of quantum operations”. *Physical Review Letters* **119**, 220502 (2017).
- [15] Nicholas J Ward, Ivan Kassal, and Alán Aspuru-Guzik. “Preparation of many-body states for quantum simulation”. *The Journal of Chemical Physics* **130**, 194105 (2009).
- [16] Hefeng Wang, S Ashhab, and Franco Nori. “Efficient quantum algorithm for preparing molecular-system-like states on a quantum computer”. *Physical Review A* **79**, 042335 (2009).
- [17] Zhang Jiang, Kevin J Sung, Kostyantyn Kechedzhi, Vadim N Smelyanskiy, and Sergio

- Boixo. “Quantum algorithms to simulate many-body physics of correlated fermions”. *Physical Review Applied* **9**, 044036 (2018).
- [18] Yordan S Yordanov, David RM Arvidsson-Shukur, and Crispin HW Barnes. “Efficient quantum circuits for quantum computational chemistry”. *Physical Review A* **102**, 062612 (2020).
- [19] Bryan T Gard, Linghua Zhu, George S Barron, Nicholas J Mayhall, Sophia E Economou, and Edwin Barnes. “Efficient symmetry-preserving state preparation circuits for the variational quantum eigensolver algorithm”. *npj Quantum Information* **6**, 1–9 (2020).
- [20] Gian-Luca R Anselmetti, David Wierichs, Christian Gogolin, and Robert M Parrish. “Local, expressive, quantum-number-preserving vqe ansätze for fermionic systems”. *New Journal of Physics* **23**, 113010 (2021).
- [21] 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**, 1–7 (2014).
- [22] Jonathan Romero, Ryan Babbush, Jarrod R McClean, Cornelius Hempel, Peter J Love, and Alán Aspuru-Guzik. “Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz”. *Quantum Science and Technology* **4**, 014008 (2018).
- [23] Harper R Grimsley, Sophia E Economou, Edwin Barnes, and Nicholas J Mayhall. “An adaptive variational algorithm for exact molecular simulations on a quantum computer”. *Nature Communications* **10**, 1–9 (2019).
- [24] Ho Lun Tang, VO Shkolnikov, George S Barron, Harper R Grimsley, Nicholas J Mayhall, Edwin Barnes, and Sophia E Economou. “qubit-adapt-vqe: An adaptive algorithm for constructing hardware-efficient ansätze on a quantum processor”. *PRX Quantum* **2**, 020310 (2021).
- [25] 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).
- [26] Panagiotis Kl Barkoutsos, Jerome F Gonthier, Igor Sokolov, Nikolaž Moll, Gian Salis, Andreas Fuhrer, Marc Ganzhorn, Daniel J Egger, Matthias Troyer, Antonio Mezzacapo, et al. “Quantum algorithms for electronic structure calculations: Particle-hole hamiltonian and optimized wave-function expansions”. *Physical Review A* **98**, 022322 (2018).
- [27] Ilya G Ryabinkin, Tzu-Ching Yen, Scott N Genin, and Artur F Izmaylov. “Qubit coupled cluster method: a systematic approach to quantum chemistry on a quantum computer”. *Journal of Chemical Theory and Computation* **14**, 6317–6326 (2018).
- [28] Yuta Matsuzawa and Yuki Kurashige. “Jastrow-type decomposition in quantum chemistry for low-depth quantum circuits”. *Journal of Chemical Theory and Computation* **16**, 944–952 (2020).
- [29] Juan Miguel Arrazola, Soran Jahangiri, Alain Delgado, Jack Ceroni, Josh Izaac, Antal Száva, Utkarsh Azad, Robert A Lang, Zeyue Niu, Olivia Di Matteo, et al. “Differentiable quantum computational chemistry with PennyLane” (2021). url: doi.org/10.48550/arXiv.2111.09967.
- [30] Scott Aaronson, Daniel Grier, and Luke Schaeffer. “The classification of reversible bit operations” (2015). url: doi.org/10.48550/arXiv.1504.05155.
- [31] Michael Reck, Anton Zeilinger, Herbert J Bernstein, and Philip Bertani. “Experimental realization of any discrete unitary operator”. *Physical Review Letters* **73**, 58 (1994).
- [32] William R Clements, Peter C Humphreys, Benjamin J Metcalf, W Steven Kolthammer, and Ian A Walmsley. “Optimal design for universal multiport interferometers”. *Optica* **3**, 1460–1465 (2016).
- [33] Hubert de Guise, Olivia Di Matteo, and Luis L Sánchez-Soto. “Simple factorization of unitary transformations”. *Physical Review A* **97**, 022328 (2018).
- [34] Mikko Möttönen, Juha J Vartiainen, Ville Bergholm, and Martti M Salomaa. “Transformation of quantum states using uniformly controlled rotations”. *Quantum Information & Computation* **5**, 467–473 (2005).
- [35] Jakob S Kottmann, Abhinav Anand, and Alán Aspuru-Guzik. “A feasible approach for automatically differentiable unitary coupled-cluster on quantum computers”. *Chemical Science* (2021).
- [36] Andrea Mari, Thomas R Bromley, and Nathan Killoran. “Estimating the gradient and higher-order derivatives on quantum hardware”. *Physical Review A* **103**, 012405 (2021).