



PAPER • OPEN ACCESS

## Exploiting subspace constraints and *ab initio* variational methods for quantum chemistry

To cite this article: Cica Gustiani *et al* 2023 *New J. Phys.* **25** 073019

View the [article online](#) for updates and enhancements.

### You may also like

- [Quantum information processing with superconducting circuits: a review](#)  
G Wendin
- [tiket: a retargetable compiler for NISQ devices](#)  
Seyon Sivarajah, Silas Dilkes, Alexander Cowtan et al.
- [Analytical framework for quantum alternating operator ansätze](#)  
Stuart Hadfield, Tad Hogg and Eleanor G Rieffel



## OPEN ACCESS

RECEIVED  
9 August 2022REVISED  
31 March 2023ACCEPTED FOR PUBLICATION  
21 June 2023PUBLISHED  
31 July 2023

Original Content from  
this work may be used  
under the terms of the  
[Creative Commons  
Attribution 4.0 licence](#).

Any further distribution  
of this work must  
maintain attribution to  
the author(s) and the title  
of the work, journal  
citation and DOI.



## PAPER

Exploiting subspace constraints and *ab initio* variational methods for quantum chemistryCica Gustiani<sup>1,\*</sup> , Richard Meister<sup>1</sup>  and Simon C Benjamin<sup>1,2,\*</sup> <sup>1</sup> Department of Materials, University of Oxford, Parks Road, Oxford OX1 3PH, United Kingdom<sup>2</sup> Quantum Motion, 9 Sterling Way, London N7 9HJ, United Kingdom

\* Authors to whom any correspondence should be addressed.

E-mail: [cica.gustiani@materials.ox.ac.uk](mailto:cica.gustiani@materials.ox.ac.uk) and [simon.benjamin@materials.ox.ac.uk](mailto:simon.benjamin@materials.ox.ac.uk)**Keywords:** quantum computing, quantum chemistry, quantum circuits, variational quantum algorithms

## Abstract

Variational methods offer a highly promising route to exploiting quantum computers for chemistry tasks. Here we employ methods described in a sister paper to the present report, entitled *exploring ab initio machine synthesis of quantum circuits*, in order to solve problems using adaptively evolving quantum circuits. Consistent with prior authors we find that this approach can outperform human-designed circuits such as the coupled-cluster or hardware-efficient ansätze, and we make comparisons for larger instances up to 14 qubits. Moreover we introduce a novel approach to constraining the circuit evolution in the physically relevant subspace, finding that this greatly improves performance and compactness of the circuits. We consider both static and dynamics properties of molecular systems. The emulation environment used is QuESTlink all resources are open source and linked from this paper.

## 1. Introduction

For over a century, the problem of simulating quantum systems has been a practical challenge in physics and chemistry. This problem is classically hard—even with many approximations—since the complexity grows exponentially in space and time. On the other hand, quantum computers are expected to be capable of tackling such tasks efficiently in a variety of scenarios. In the current phase of quantum technology, the so-called NISQ [1] era (*noisy intermediate-scale quantum*), we endeavour to find the first practical examples of such a quantum advantage.

Having a NISQ device means having access to hundreds to thousands of qubits. This scale is, however, not enough to run fully fault-tolerant computations, thus limits us to run only shallow circuits. Variational quantum algorithms (VQAs) are the most widely studied class of algorithms that are potentially compatible with NISQ device constraints. The VQAs operate on parameterised quantum circuits (ansätze) and rely on classical processors for the parameter evolution, therefore, the quantum circuits remain shallow. The shared quantum/classical nature of the paradigm means that these approaches are also called ‘hybrid’ quantum algorithms.

An eigensolver—finding the ground or other eigenstates of a Hamiltonian of interest—was the first VQA that was proposed nearly a decade ago [2]. The ground-state energy problem is provably hard, known for having complexity QMA-complete [3]. Since then, many VQAs have emerged for various applications, *e.g.*, chemistry [4–12], solving equations [13–16], and quantum circuit compilations [7, 17–20]. We refer to [12, 21–23] for overviews and current progress reports on the VQAs.

The choice of ansatz is critical for a successful optimisation, where it must be expressive enough to approximate the ground state, but not contain too many parameters as it may cause vanishing gradients [24]. Several strategies have been proposed to devise a good ansatz, *e.g.*, problem-inspired ansätze [25, 26], hardware-efficient ansätze [27–30], unitary coupled cluster ansätze [4, 8, 30, 31], and adaptive structure ansätze [4, 29, 30, 32]. We refer to [21, 22] for an overview of ansatz constructions.

Our focus in this work is on techniques where ansätze are evolved on the fly by random guesses followed by the application of selection criteria and judicious deletion. We only consider noise-free gates, and we performed our algorithms on an emulated quantum computer with up to 14 qubits. This work is accompanied by our study on various existing and new optimisation technicalities in the sister paper [33]; here, we apply techniques we found to perform well.

We test our VQA techniques to solve well-known chemistry problems: finding the ground state of a Hamiltonian and synthesising shallow circuits which implement the time-evolution operator. For the latter, we consider both full Hilbert space dynamics, as well as circuits created under the more permissive condition that dynamics should be correct in a specific subspace; specifically, this is in a subspace that preserves the occupation number. The subspace constrained time-evolution operator (propagator) is obtained via the *subspace compilation technique* that is introduced in our sister paper [33]. Our interest in the subspace-restricted propagators stems from the fact that most chemical reactions preserve the total number of particles, making processes outside of this subspace irrelevant. We find that it is more efficient to employ propagators that operate in a subspace rather than in the entire Hilbert space. To our knowledge, no prior method to construct such propagators has been explored in the literature.

Our technique is agnostic to the Hamiltonian and can operate without or with only trivial knowledge of the system. We test it on finding the ground state of small molecules, using up to 14 qubits. We obtained remarkable gate count reductions compared to other techniques that we are aware of. For instance, we only need a total of 20 gates to approximate the ground state of the LiH molecule, which is a 12-qubit problem. We also test our technique to synthesise the time-evolution operators of the H<sub>2</sub> molecule dynamics over various time intervals. Interestingly, our circuits have nearly-constant gate counts: no more than 18 gates for the subspace restricted compilation and no more than 34 gates for the full space compilation; this is considerably fewer than the canonical *Trotter-Suzuki* product formulas that need hundreds of gates, even with the optimal choice of Trotter order. Our shallow propagators can be concatenated and thus potentially pave the way to achieving longer periods of modelled time even in NISQ simulations.

The layout of this paper is the following. Firstly, in section 2, we discuss the general formalism and the optimisation frameworks: section 2.1 for the ground state estimation problem and section 2.2 for the circuit synthesis of the time-evolution operator constrained to an arbitrary subspace. Secondly, section 3 outlines the optimisation techniques. Then, section 4 shows all the numeric results together with discussions and analysis. Section 4 is divided into two subsections by the problems: section 4.1 for the ground state estimation and section 4.2 for the circuit synthesis of the time-evolution operator. Finally, section 5 concludes this work followed by a general discussion.

## 2. Chemistry problems with 2nd quantised Hamiltonians

Estimating and preparing the ground state and implementation of the time-evolution operator are the first steps toward dynamical molecular simulation. Here, we tackle these problems by means of VQAs.

To simulate a molecular Hamiltonian on a quantum computer, one needs to represent the problem in a format that can be solved on a quantum computer, e.g., by operating quantum gates. In our study, we describe the electronic Hamiltonian using second quantization, namely


$$H = \sum_{i,j} h_{ij} a_i^\dagger a_j + \sum_{i,j,k,l} h_{ijkl} a_i^\dagger a_j^\dagger a_k a_l, \quad (1)$$

where  $\{a_j\}$  and  $\{a_j^\dagger\}$  are lowering and raising operators acting on a Fock space. The coefficients  $h_{ij}$  and  $h_{ijkl}$  are overlaps and exchange integrals on the specified basis set; these can be efficiently computed on classical computers. We use the basis set STO-3G [34], namely, contracted three Gaussian functions in which parameters are obtained by a self-consistent field procedure.

We proceed by mapping the Fock space onto the Hilbert space of the qubits. We use the canonical Jordan–Wigner transformation [35, 36] to map the Fock states into qubit states in the computational basis and the Fock operators into quantum gates. Finally, our Hamiltonian is expressed in the form of a Pauli sum,

$$H = \sum_k c_k P_k, \quad (2)$$

where  $c_k \in \mathbb{R}$  and  $P_k \in \{I, \sigma_x, \sigma_y, \sigma_z\}^{\otimes n}$  are  $n$ -qubit Pauli strings. We utilise the Python packages *Openfermion* [37] and *PyScf* [38] to generate our Hamiltonians. After the transformations, our electronic problems are ready to be processed on a quantum computer. From this point onwards, we assume the Hamiltonians have the form of equation (2).



**Figure 1.** Naive Trotterisation on an  $H_2$  molecule's propagator for  $\Delta t = 2$  up to the fourth order. Each order show Trotterisations with various trotter number ( $n$ ), consistently increasing from  $n = 1$  up to  $n = 56$  for the first order; thus  $n$  coincides with the number of markers. The error is defined as the matrix distance (equation (11)) between the propagator and the Trotterised circuit. We refer to appendix A for details on the Trotter formula, the circuits, and circuit simplification.

## 2.1. Finding the ground state

VQAs discover the ground state of a Hamiltonian  $H$  by minimising the energy expectation value (cost function)

$$\min_{\theta} \langle \psi_0 | C^\dagger(\theta) H C(\theta) | \psi_0 \rangle, \quad (3)$$

where  $C(\theta)$  is a parameterised quantum circuit (ansatz) and  $|\psi_0\rangle$  is an initial fixed state. By linearity, the cost function can be written as

$$\sum_k c_k \langle \psi_0 | C^\dagger(\theta) P_k C(\theta) | \psi_0 \rangle. \quad (4)$$

Quantum computers can efficiently compute each term above independently, therefore, the circuits can remain shallow.

When the iterations reach convergence by some criterion, *e.g.*, cost improvement from the previous iteration is less than  $\delta$ , the optimisation halts and returns the optimised parameters  $\theta'$ . Then, the resulting ground state energy estimation is  $E_{gs} \approx \langle \psi_0 | C(\theta')^\dagger H C(\theta') | \psi_0 \rangle$ , with the wave function  $C(\theta') | \psi_0 \rangle$ , and has an error  $\varepsilon$ . The error  $\varepsilon$  is defined as the difference to the exact ground state energy. Our goal is for this final energy estimate error  $\varepsilon$  to be within the chemical accuracy  $\varepsilon = 1.59 \times 10^{-3}$  Hartree (1 kcal/mole), which is the typical error relevant to thermodynamic experiments [39].

## 2.2. The time-evolution operator in the subspace

The time-evolution operators,


$$U(\Delta t) = e^{-i\Delta t H}, \quad (5)$$

are usually hard to compute for interesting problems, *i.e.*, when Hamiltonian terms are not mutually commuting. The canonical way to obtain  $U(\Delta t)$  is by employing the Suzuki-Trotter formula [40]; this process is called *Trotterisation*.

The Suzuki-Trotter formula approximates  $U(\Delta t)$  by the products of the exponents, in which the error term is improved by slicing the total time  $\Delta t$  into  $n$  uniform subintervals, also known as *Trotter number*. Increasing  $n$  results in a larger gate count, trading off with an improved error (equation (11)), as demonstrated in figure 1. It is worth mentioning that, the gate count can be improved with more rigorous gate simplification by taking into account the ordering, which we do not consider here. Nevertheless, by using naive Trotterisation and simple gate removals/simplifications, from figure 1, we need around 900 gates to obtain an error of 0.001, which is quite immense for a 4 qubit problem!

In this paper, we demonstrate our circuit synthesis techniques by creating compact circuits for  $H_2$  propagators; we do so for various durations  $\Delta t$ . We take advantage of the block diagonal structure of propagators to synthesise very shallow circuits ( $\leq 18$  gates) constrained to a subspace, with a target error of less than 0.001 for various  $\Delta t$ . Our sister paper [33] shows that circuit synthesis in a subspace is more efficient than in the entire space, given that the unitary has a block diagonal form. In the following, we summarise the idea of the subspace compilation technique; we refer to our sister paper [33] for further details.

Given a propagator  $U(\Delta t)$ , we synthesise a hardware-efficient circuit that approximates it by means of VQA. Ideally, the operator  $U(\Delta t)$  can be implemented on the quantum computer with a great accuracy. For instance,  $U(\Delta t)$  is a high order Trotterisation with a very deep circuit. Our VQA is aimed at discovering a



**Figure 2.** The circuit setup for computing the cost function. A maximally entangled state  $|\Phi\rangle_{RA}$  (equation (7)) is prepared in registers  $R$ - $A$  by applying a preparation operator  $P$  from a zero state  $|0\rangle^{\otimes(d+m)}$ . This technique is a generalisation of the circuit compilation introduced in [41]. The propagator  $U(\Delta t)$  is applied to the first register; it is the unitary target, e.g., a long canonical circuit obtained by Trotterisation. Then, a parameterised ansatz  $\mathcal{C}(\theta)$  is applied, followed by measurement in basis  $|\Phi\rangle_{RA}$ . The measurement outcomes are accounted for in the cost function.

much shorter circuit that approximates  $U(\Delta t)$  in the full Hilbert space and in a subspace. The key to our circuit synthesis is evaluation of the cost function shown in figure 2.

Let  $\mathcal{H}_S$  be the subspace of  $\mathcal{H}$  that preserves the relevant occupation number; this method requires  $m$  ancillae, where  $m = \lceil \log_2(\dim \mathcal{H}_S) \rceil$ . By referring to figure 2, the core of our circuit synthesis technique is the minimization of

$$\min_{\theta} (1 - \langle \Phi |_{RA} U(\Delta t) \mathcal{C}(\theta) | \psi \rangle_{RA}), \quad (6)$$

where  $|\Phi\rangle_{RA}$  is maximally entangled state between spaces  $R$  and  $A$ , namely,

$$|\Phi\rangle_{RA} = \frac{1}{\sqrt{2^m}} \sum_{s_j \in S} |s_j\rangle_R |j\rangle_A. \quad (7)$$

It means  $\text{Tr}_A(|\Phi\rangle\langle\Phi|)_{RA} = I_R$  and  $\text{Tr}_R(|\Phi\rangle\langle\Phi|)_{RA} = I_A$ , where  $I$  is an identity matrix. Here, the basis  $\{|s_j\rangle_R\}$  spans  $\mathcal{H}_S$ , which is the subspace of  $\mathcal{H}$  spanned by basis  $\{|s_j\rangle_{RA}\}$ . Upon convergence of the cost function (equation (6)) over a set of parameters  $\theta$ , the time evolution operator is approximated by  $U(\Delta t) \approx^\epsilon \mathcal{C}(\theta)^\dagger$  with some error  $\epsilon$ . We define the error  $\epsilon$  as a matrix distance between  $U(\Delta t)$  and  $\mathcal{C}(\theta)^\dagger$ ; the definition of  $\epsilon$  is elaborated later in section 4.2.

### 3. Discussion on the choice of optimisation

The performance of a VQA scheme crucially depends on the setup of the optimiser. We outline the core aspects of our optimization techniques below, which are elaborated in greater detail in our sister paper [33], to which we refer the interested reader.

#### 3.1. Gate pool

We use a large set of gates comprising parameterised single-qubit and controlled rotations assuming full connectivity, namely  $\{R\sigma_i(\theta), C_i(R\sigma_j(\theta))\}_{\sigma \in \{x,y,z\}}$ , where  $i, j \in \{0, \dots, n-1\}$  and  $i \neq j$ , for a  $n$ -qubit problem. The parameter sets are dense,  $\theta \in \mathbb{R}$ , since most experiments allow continuous parameters. Notice that the number of parameters coincides with the number of gates.

#### 3.2. Adaptive ansatz

The ansatz is constructed by trying out several candidates of random ansätze for multiple rounds. On each round, every ansatz candidate is initialized by a random circuit, comprising legitimate random gates, which are in turn initialized with random parameters to be optimised. The best candidate is chosen by the lowest cost function, which then continues to evolve through the optimisation. This technique is inspired by the ADAPT technique [4], in which the global minimum is approached with small steps along the evolving ansatz. For the ground state minimization, the initial state is a product state with the correct occupation number; our choice is  $|\psi_0\rangle = |0 \dots 001 \dots 11\rangle$ , by least significant bit convention, where the number of ones coincides with the occupation number. We refer to the *random search* technique in [33] for further details.

#### 3.3. Removal of superfluous gates

At each step of ansatz evolution, judicious deletions are performed on the ansatz before evolving further. This step is crucial, since superfluous gates can trap the system in the local minima, and a large number of

optimisation parameters can induce a Barren plateau. Our sister paper [33] elaborates on this removal technique in section IIID.

### 3.4. Parameter optimisation

We employ the natural gradient descent (or imaginary-time evolution) method [42] to train our parameters. The free parameters  $\theta$  are updated as follows

$$\theta' = \theta - \lambda F(\theta)^{-1} \nabla C(\theta), \quad (8)$$

where  $\lambda$  is the gradient step,  $\nabla C(\theta)$  is gradient of the cost, and

$$F_{ij} = 4\text{Re} [\langle \partial_i \psi(\theta) | \partial_j \psi(\theta) \rangle - \langle \partial_i \psi(\theta) | \psi(\theta) \rangle \langle \psi(\theta) | \partial_j \psi(\theta) \rangle] \quad (9)$$

is the *quantum Fisher information metric* [43, 44], also known as *quantum metric tensor*. Note that the equation above holds only for pure states. The metric tensor gives additional information in the state space from perturbation in the parameter space; this improves the optimisation performance as demonstrated in [42]. Also, to speed up optimisation, we exponentially increase lambda within a single iteration until we arrive at a local minimum along the natural gradient direction  $F^{-1}(\theta) \nabla C(\theta)$ .

### 3.5. Hyperparameters

Several hyperparameters are set according to the problem, usually by the number of qubits and estimated difficulty. These include the number of random gates inserted at once, the proposition of one- and two-qubit gates insertion, convergence criteria, the gradient step ( $\lambda$  in equation (9)), the regularization parameter for  $F^{-1}(\theta)$  (equation (8)), and some parameters relevant to the gates removal subroutines. Moreover, some hyperparameters are updated when the optimisation fails to lower the cost function.

## 4. Results and analysis

The results reported here are accessible in our GitHub repository see supplementary material. It is important to note that we do not consider noise in our simulations. We use QuESTLink [45] to emulate our algorithms, which is the Mathematica interface to a compiled form of the QuEST [46] toolset (a highly efficient c codebase). We employ a GPU to accelerate our computations. As Python is popular in the relevant communities, it is worth mentioning that one can also reproduce our simulations using pyQuEST [47], which is the Python interface to QuEST.

### 4.1. Ground state estimations

We test our optimisation technique to estimate the ground states of small molecules that can be represented by up to 14 qubits. We use the stable geometry for the molecules and constant distance for the (artificial) hydrogen chains. The geometry of each molecule is shown in figure 3.


We terminate the optimization when the updated energy fails to lower the cost function by more than  $\delta = 10^{-6}$ . Each molecule is run for ten trials. Our best results—by the total gate count—among the trials, within chemical accuracy, are shown in figure 4.

The resulting gate counts are remarkably low for some small molecules, such as  $\text{H}_2$ ,  $\text{He}_2\text{H}^+$ , and  $\text{He}_3\text{H}^+$ . The initial states are already somewhat close to the ground state in these cases. We note that, for the  $\text{He}_2\text{H}^+$  case, the initial state can easily be transformed into the ground state: we start with state  $|011111\rangle$ , where the ground state is  $0.071669|011111\rangle + 0.997428|111101\rangle$ . That is achievable with a hadamard and a controlled-rotation gates. Our optimiser discovered circuit  $Rx_5(-0.1453)C_5[Rx_1(\pi)]$  to approximate the ground state of  $\text{He}_2\text{H}^+$  with an energy error magnitude of  $10^{-10}$ .


To put our results into context, we compare them with the results of various methods tabulated in table 1. We emphasise that this comparison is not straightforward due to a vast landscape of conventions. We observe differences in gate pools, unreported total gate count, objective functions, accuracies, molecular geometries, and Hamiltonian mappings may lead to unfair comparisons. For instance, the results from MoG-VQE [32] show small numbers of CNOT gates but large numbers of single-qubit gates. This is no doubt because the MoG-VQE method focuses on reducing CNOT gates while single-qubit gates are regarded as plentiful and placed without simplification. Nonetheless, in most cases, we find lower numbers of total gates compared to other results even when we decompose our circuits into CNOT and single rotation gates. For instance, figure 5 shows our circuit to prepare the ground state of LiH, which comprises only 20 gates.

Although our results show small gate counts, our ansatz choice comes with a drawback due to the growing number of parameters. For instance, in constructing the ansatz, sometimes, the appropriate gates





**Figure 3.** Geometry of each molecule used in this work. The units are displayed in Angstrom ( $\text{\AA}$ ).



(a)

Mole- cule	Pauli terms	Two- qubit gate	Single- qubit gate
H <sub>2</sub>	15	3	1
H <sub>3</sub>	62	8	3
H <sub>4</sub>	185	36	10
He <sub>2</sub> H <sup>+</sup>	62	1	1
He <sub>3</sub> H <sup>+</sup>	193	1	2
LiH	631	12	8
BeH <sub>2</sub>	666	39	19
H <sub>2</sub> O	1086	49	37

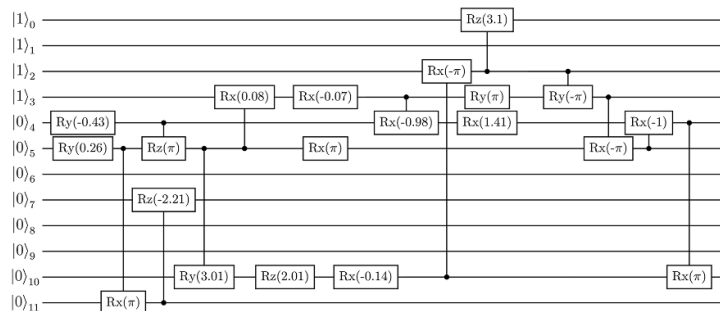
(b)

**Figure 4.** The gate count obtained by our optimiser to prepare the ground state within chemical accuracy. The number of parameters coincides with the gate count. (a) The gate count plot versus the number of qubits. (b) The gate count for two- and one-qubit gates compared to the number of Pauli terms in the Hamiltonian.

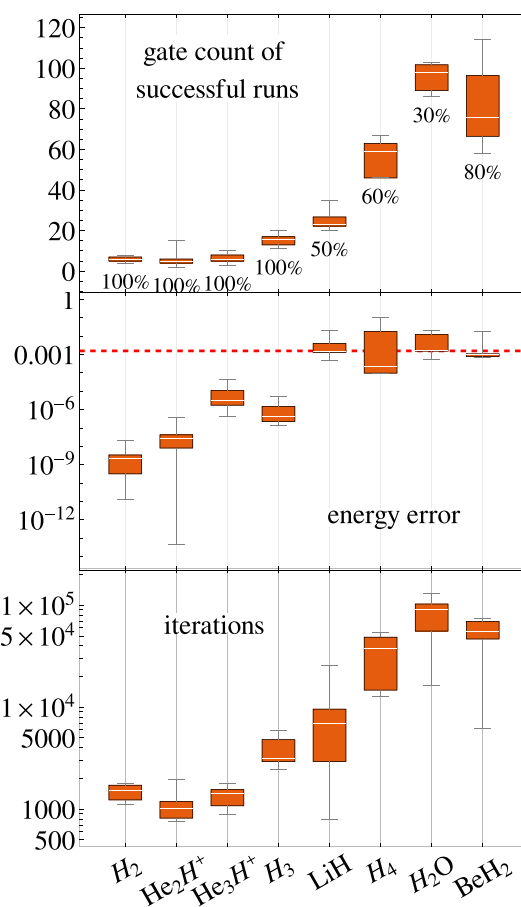
**Table 1.** The gate and parameter counts from various methods. Methodological differences to ours are summarised in the description column. We select these results because of similarity in the molecular geometry and data availability.

Mole-cule	Gate count	Param count	Description
H <sub>2</sub>	125	2	ADAPT-VQE [48].
	105	70	Fixed ansatz [49], with 35 CNOT and 70 z-rotation gates, has 8 qubits, with STO-6G basis, and nearest-neighbour connectivity.
H <sub>4</sub>	>150	26	AVQITE [50], based CNOT count with cyclic nearest-neighbour connectivity, and tapered to 6 qubits.
	279	248	MoG-VQE [32], with bond distance 1.2 Å, 31 CNOT and 248 single rotations gates.
	>2208	11	qubit-ADAPT [30], based on CNOT count, and within accuracy $10^{-14}$ .
LiH	108	72	Fixed ansatz [49], with 36 CNOT and 72 z-rotation gates, tapered to 10 qubits, and has nearest-neighbour connectivity.
	108	96	MoG-VQE [32], with 12 CNOT and 96 single rotation gates.
	>6824	30	qubit-ADAPT [30], based on CNOT count, and within accuracy $10^{-10}$ .
BeH <sub>2</sub>	>400	46	AVQITE [50], based on CNOT count with cyclic nearest-neighbour connectivity, and tapered to 12 qubits.
	81	72	MoG-VQE [32], with 9 CNOT and 72 single rotation gates, and tapered to 8 qubits.
H <sub>2</sub> O	>1550	84	4-UpCCGSD-PECT ansatz [31], and count based on circuit depth.
	>100	5	AVQITE [50], based on CNOT count with cyclic nearest-neighbour connectivity, and tapered to 8 qubits.

are discovered by creating a large random circuit followed by judicious deletions. Since large circuits come with many parameters, our method might be prone to a barren plateau on larger scales. On the other hand, the techniques that employ coupled-cluster family operators such as qubit-ADAPT and ADAPT-VQE show large gate counts but have a small number of parameters; these techniques might be more scalable than ours.




**Figure 5.** The LiH ground state preparation circuit synthesised by our algorithm. The LiH molecule is commonly used to benchmark VQE algorithms. The circuit is rounded to the multiple of  $\pi$  or to two-decimal places; it has energy error of  $\varepsilon = 1.4 \times 10^{-3}$  Hartree. Three qubits remain untouched; two of them could have been manually identified as redundant by tapering them off [51], *i.e.*, reducing the number of qubits by preserving some symmetries in the Hamiltonian.



**Figure 6.** Optimisation results from all trials for each molecule. The *gate count* section shows the successful run defined by energy errors within chemical accuracy. The *success rate*, *i.e.*, percentage of trials that converge below the threshold  $\delta = 10^{-6}$ , are shown below each whisker. The water molecule shows the lowest rate of successful runs. In the *energy error* section, the dashed red line indicates chemical accuracy. The large molecules show more results converged to the local minima. Finally, the *iterations* section shows the total gradient evaluations, including trying all candidates and judicious deletions.

In the following, we discuss the factors influencing the difficulty of preparing the ground state and how they affect the performance of our optimisations. First, we know that the difficulty in discovering the ground state increases with the size of the molecule and the number of terms in the Hamiltonians, but, they are not the only factors. For instance, the ground state of  $\text{He}_3\text{H}^+$ , which is represented using eight qubits, requires only three gates; in contrast to  $\text{H}_3$ , which has six qubits but requires 11 gates. Also, for an equally big 8-qubit problem,  $\text{H}_4$  needs 46 gates accompanied with a lower success rate and needs more iterations—see figure 6, even though it has fewer Pauli terms. Since the task of estimating the ground state problem is proven to be





QMA-complete [3], rigorously quantifying the optimisation difficulty of a Hamiltonian of practical scale may remain an unfeasibly hard challenge.

Intuitively, if we knew that the ground state was highly correlated, we would expect the optimisation to be difficult. A highly correlated state usually corresponds to a highly entangled state, which requires a sophisticated circuit with numerous entangling gates to produce. Thus, circuits preparing highly entangled ground states are more difficult to synthesise as they need deep circuits and (or) large number of parameters which are likely to yield flat gradient manifolds. Therefore, in the following, we quantify the entanglement of the target state and evaluate its impact on the performance of our optimiser.

Having the advantage of an even number of qubits in our molecules, we can straightforwardly use the bipartite entanglement measure between two chosen subsystems. A simple yet powerful measure for pure states is the *Schmidt measure* (*Hartley strength*) [52]:

$$E(|\psi\rangle) = \log_2(\text{rank}(|\psi_{\text{sch}}\rangle)), \quad (10)$$

where  $|\psi_{\text{sch}}\rangle$  is the Schmidt form of  $|\psi\rangle$ . The unit of Hartley strength is known as ‘ebit’, where Bell states correspond to one ebit.

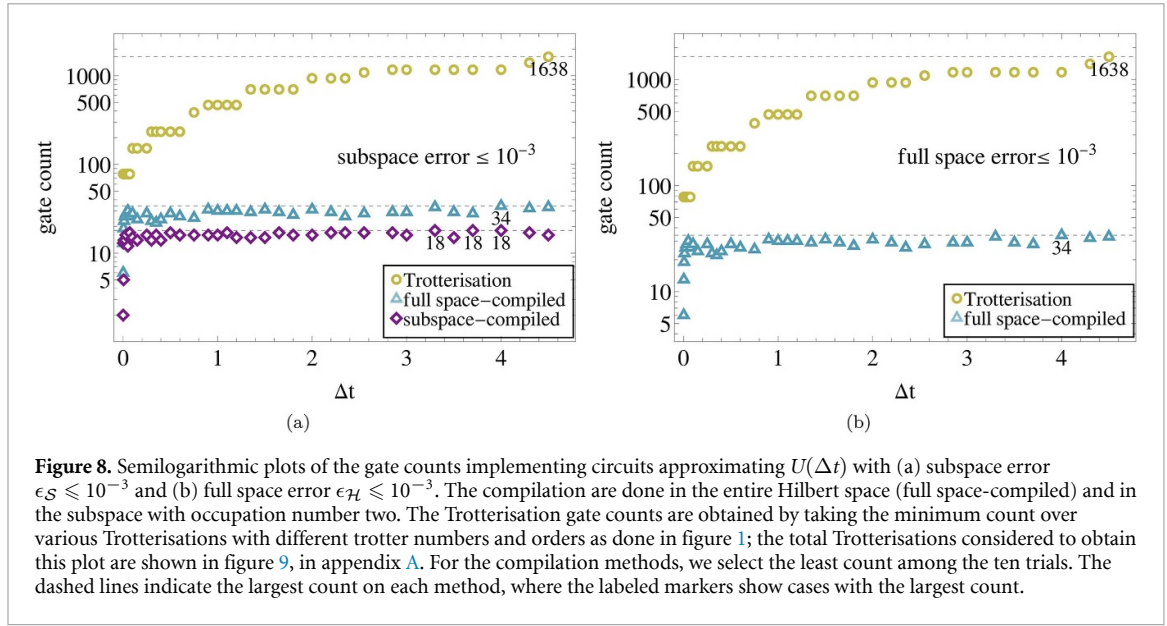
Indeed, as shown figure 7, the entanglement measure shows a stronger correlation to the entangling gate count than the number of qubits or the Pauli terms. For instance, the two-qubit gate counts of the molecules  $H_2$ ,  $He_2H^+$ , and  $He_3H^+$  are almost identical for the same Hartley strength. Consistent results also shown for  $H_3$  and  $LiH$ , although  $LiH$  was more difficult for us to optimise, with only a 50% success rate (the error within chemical accuracy). In addition,  $H_4$  is expected to be more difficult than  $LiH$  and  $H_3$ , which is also reflected in the optimisation performance. However, in the case of  $H_2O$  and  $BeH_2$ , our optimiser found that  $H_2O$  is harder than  $BeH_2$ . The performance of our optimiser relies on hyperparameters which are manually set. It might be that the hyperparameters we chose for  $H_2O$  are not optimal for this case.

Our optimiser has successfully learned shallow circuits that prepare the ground states with gate counts much lower than  $\Theta(n^3 4^n)$ , which is a bound on the number of two-qubit gates necessary to prepare an arbitrary  $n$ -dimensional unitary gates [53]. Moreover, some of our results outperform results by other methods as shown in table 1. However, the performance of our optimiser relies on randomness and a good choice of hyperparameters. For instance, in some cases, we insert a hundred random gates in which 99 of them are deleted and only one remains; this is a significant overhead in guessing a single gate. Many trials are sometimes required to guess the proper hyperparameters as well. Therefore, we do anticipate the need for adaptations and refinements in scaling the method to much larger molecules.

#### 4.2. The time evolution operator circuits in the subspace and full space

Using the same optimisation technique, we synthesise circuits that implement time evolution operators of the  $H_2$  molecule with various time steps  $\Delta t$ . The optimisations are performed in the entire Hilbert space and in the subspace that preserves the correct number of particles. Interestingly, as we will see later, the resulting gate counts remain relatively constant while increasing  $\Delta t$ .

First, we formalise the operator that quantifies the quality of the synthesised circuits. It is natural to compare two matrices by a distance norm to quantify accuracy. By referring to figure 2, the error of an approximation  $U(\Delta t) \approx C^\dagger(\theta)$  that is obtained by recompilation in a subspace  $S \subseteq \mathcal{H}$  is defined as



**Figure 8.** Semilogarithmic plots of the gate counts implementing circuits approximating  $U(\Delta t)$  with (a) subspace error  $\epsilon_S \leq 10^{-3}$  and (b) full space error  $\epsilon_H \leq 10^{-3}$ . The compilation are done in the entire Hilbert space (full space-compiled) and in the subspace with occupation number two. The Trotterisation gate counts are obtained by taking the minimum count over various Trotterisations with different trotter numbers and orders as done in figure 1; the total Trotterisations considered to obtain this plot are shown in figure 9, in appendix A. For the compilation methods, we select the least count among the ten trials. The dashed lines indicate the largest count on each method, where the labeled markers show cases with the largest count.

$$\epsilon = \min_{\phi} \|\Pi_S (U(\Delta t) - e^{i\phi} \mathcal{C}(\theta)^\dagger) \Pi_S\|_{\infty}, \quad (11)$$

where  $\Pi_S$  is the projector to the subspace  $\mathcal{S}$ , and  $\|\cdot\|_{\infty}$  is Schatten- $\infty$  norm operator defined as

$$\|A\|_{\infty} = \max\{\|Av\| : v \in \mathcal{S}\}, \quad (12)$$

which can be obtained by taking the largest magnitude eigenvalue. For practicality, we say ‘subspace error’,  $\epsilon_S$ , if defined over  $\mathcal{S}$  and ‘full space error’,  $\epsilon_H$ , if defined over  $\mathcal{H}$ . In the context of circuit synthesis, we interchangeably use error and distance for clarity.

We perform optimisations in the subspace  $\mathcal{S}$  and in the full space  $\mathcal{H}$ , where  $\mathcal{S}$  is the subspace with occupation number two—as present in  $H_2$  molecule; thus,  $\dim(\mathcal{S}) = 6$  and  $\dim(\mathcal{H}) = 16$ . We set the cost convergence condition to  $\delta \leq 10^{-7}$ , where the cost is defined in equation (6). For such a  $\delta$ , we obtain a distance in of the order  $10^{-4}$ . We synthesise propagator circuits for 37 different  $\Delta t$  values, ranging from 0.001 to 4.5. For each  $\Delta t$ , we try ten independent calculations. The best results of all trials with errors  $\epsilon_H, \epsilon_S \leq 10^{-3}$  are shown in figure 8. In that figure, we also compare our results with the Trotterisation method.

Interestingly, figure 8 shows that the circuits synthesised by our optimiser have a relatively constant number of gates, for both the compilation in the subspace and the full space. Moreover, the gate counts we have obtained are very low: no more than 18 for the subspace compilations and no more than 34 for the full space compilations; those are significantly smaller than the gate counts we have obtained via Trotterisation, which grow exponentially with  $\Delta t$ .

Figure 8(a) shows that subspace-compiled gate counts are always smaller than the full space-compiled ones. This is consistent with an observation in our sister paper [33]: the subspace compilation is more efficient and performs better than full space compilation. One can see it as reducing the problem size by compiling in the subspace. We provide details on the performance of our circuit synthesis in appendix B.

In this numerical experiment, we observe that the complexity of the unitary is practically the same for different  $\Delta t$  with our technique. This may be due to the nature of parameterised ansatz optimisation that trains the parameters within a particular unitary structure. For instance, this technique was tested on a dense random unitary in our sister paper, resulting in significantly more complex optimisations and showing much larger gate counts. Fortunately, the time-evolution operators have a constant unitary structure for any  $\Delta t$ . Its sparse structure and block diagonal form are ideal to our compilation techniques.

Our results may provide a significant advantage in some practical applications by the following schemes. First, a long-duration time-evolution simulation by repeated application of a propagator is more feasible due to the shallow circuit of each iterate. Second, one can compile a long-duration time-evolution propagator by recursive compilation to obtain a short circuit to simulate a longer duration. The advantage is even more significant for reactions that preserve the number of particles. Finally, it possesses a potential in first-quantised simulations [54–56]. For instance, Chan *et al* [56] have successfully predicted oscillations in

the Helium dimer using the first-quantised method. Their simulation extensively employs propagators for each potential term, spatial grid, and time grid; thus, maintaining shallow circuits of the propagators can significantly improve the efficiency, which are expected to assist in the prediction of the behaviour of larger molecules in the future.

## 5. Conclusion and further discussion

We implemented optimisation techniques studied in our sister paper [33]: (1) ground state estimations of molecules  $\text{H}_2$ ,  $\text{H}_3$ ,  $\text{H}_4$ ,  $\text{He}_2\text{H}^+$ ,  $\text{He}_3\text{H}^+$ ,  $\text{LiH}$ ,  $\text{BeH}_2$ , and  $\text{H}_2\text{O}$ , and (2) synthesising circuits of the  $\text{H}_2$  molecule time-evolution operators for various time durations  $\Delta t$  in which optimisations are run in the subspace that preserves the number of particles, as well as in the entire space. Key results from our study here are the following. For problem (1), we obtained shallow circuits to prepare the ground state of each molecule summarised in figure 4. For problem (2), we synthesised shallow circuits with a practically constant depth for all given  $\Delta t$ , which is summarised in figure 8. We found that compilation in the subspace reduces the gate count versus the full space compilations by a factor of about 2.

We can anticipate some of the challenges associated with scaling the technique to problems beyond the quantum advantage threshold. First, it might not scale well since it relies on randomness to find gates that can lower the cost, and the parameter count grows with the number of gates simultaneously inserted. This would likely necessitate a more nuanced protocol for gate introduction in place of the present purely-random heuristic. Second, the performance relies on the hyperparameters that are chosen manually. The solution, automatic selection of hyperparameters, would be an interesting machine learning task in its own right. Regardless of these caveats, our optimiser has already demonstrated potential to solve small molecule chemistry problems with remarkable efficiency.

As an additional note, it was challenging for us to comprehensively characterize the outcomes of our method with the literature due to the lack of a common framework and transparency. For instance, in our opinion, the total gate count is as important as the two-qubit gate count and the circuit depth for the following reasons. First, by the state of the art of quantum computing, *i.e.*, in the NISQ era, most experiments do not yet exploit the parallelism of quantum gates. Second, performing CNOT gate and arbitrary controlled- $x$  rotations without fault-tolerance is equally difficult in experiments. Finally, VQE is aimed at non-fault-tolerant devices such as NISQ devices. Hopefully, the upcoming works in this area will show more transparent results and consider the experimental situations to push the practical applications of NISQ forward.

## Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://github.com/cicacica/chemistry-code>.

## Acknowledgments

We thank Daniel Marti Dafcik, Hans Hon Sang Chan, and Bálint Koczor for the discussions on the variational algorithms and chemistry problems. We thank Tyson Jones for developing QuESTlink and writing functionalities such that we can conveniently implement our ideas. C G thanks Jonathan Conrad for the proof reading. C G and S C B acknowledge financial support from EPSRC Hub grants under Agreement No. EP/T001062/1, from the IARPA funded LogiQ project, the EU flagship AQTION project, European Union's Horizon 2020 research, and innovation programme under Grant Agreement No. 951852 (QLSI).

## Appendix A. Trotterisation

Given that the Hamiltonian can be represented as

$$\sum_k c_k P_k, \quad \text{where } c_k \in \mathbb{R}, P_k \in \{X, Y, Z, I\}^{\otimes n}.$$

We partition all of the Pauli terms  $\{P_k\}$  into subsets  $A$  and  $B$  in which elements are commuting to each other. The following equations are the first fourth orders of Suzuki–Trotter formulas [36, 40] used in this paper:

$$e^{(A+B)t} \approx (e^{At/n} e^{Bt/n})^n + O(t^2) \quad (A1)$$


$$e^{(A+B)t} \approx \underbrace{(e^{At/2n} e^{Bt/n} e^{At/2n})(e^{Bt/2n} e^{At/n} e^{Bt/2n})(e^{At/2n} e^{Bt/n} e^{At/2n}) \dots}_{n \text{ terms}} + O(t^3) \quad (A2)$$

$$e^{(A+B)t} \approx (e^{\frac{7}{4}At/n} e^{\frac{2}{3}Bt/n} e^{\frac{3}{4}At/n} e^{\frac{-2}{3}Bt/n} e^{\frac{-1}{24}At/n} e^{Bt/n})^n + O(t^4) \quad (A3)$$

$$e^{(A+B)t} \approx \left( \prod_{i=1}^5 e^{p_i At/2n} e^{p_i Bt/n} e^{p_i At/2n} \right)^n + O(t^5), \quad p_1 = p_2 = p_4 = p_5 = \frac{1}{4 - 4^{1/3}}, p_3 = 1 - 4p_1. \quad (A4)$$

We use a slightly different ordering for the second order Trotter in equation (A2) because this ordering shows better error than the canonical one.

We use the canonical exponentiating circuit form, for example,



The circuit is then simplified by the following operations:

$$\begin{aligned} CNOT_{i,j} CNOT_{i,j} &= I, \\ Rx_i(a) Rx_i(b) &= Rx_i(a+b), \\ Rx_i(+a) Rx_i(-a) &= I, \\ H_i H_i &= I. \end{aligned}$$

Note that we do not consider commutations in simplifying the gates, which can improve the gate counts.



We generate Trotter circuits for various durations,

$$\Delta t \in \{0.001, 0.002, 0.005, 0.007, 0.01, 0.02, 0.03, 0.035, 0.05, 0.07, 0.1, 0.15, 0.25, 0.3, 0.35, 0.4, 0.5, 0.6, 0.75, 0.9, 1, 1.1, 1.2, 1.35, 1.5, 1.65, 1.8, 2, 2.2, 2.35, 2.55, 2.85, 3, 3.3, 3.5, 3.7, 4, 4.3, 4.5\}. \quad (A5)$$

For every  $\Delta t$ , we find Trotter circuits for the first fourth order with various Trotter numbers  $n$ , where

$$\begin{aligned} n &\in \{1, \dots, 56\} && \text{for the first order,} \\ n &\in \{1, \dots, 36\} && \text{for the second order,} \\ n &\in \{1, \dots, 18\} && \text{for the third order, and} \\ n &\in \{1, \dots, 5\} && \text{for the fourth order.} \end{aligned} \quad (A6)$$

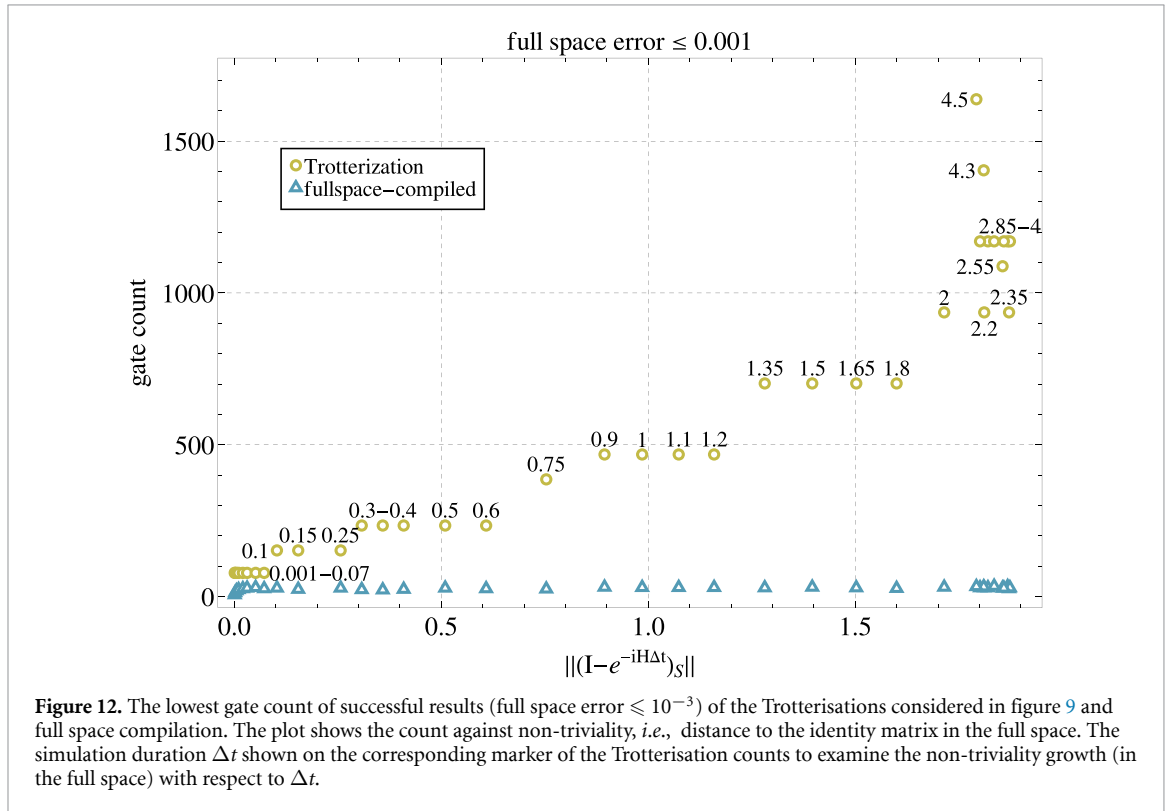

For each Trotter circuit, we compute the error for the full space and subspace as shown in figure 9. The error is defined by the matrix distance to the exact matrix that is obtained by Mathematica function `MatrixExp`.



## Appendix B. Circuit synthesis of the time-evolution operator

We synthesise circuits of the time-evolution operator of  $H_2$  molecule for various  $\Delta t$  defined in equation (A5). To justify that we are not running some trivial problems, we compute ‘non-triviality’, defined as the distance — for both subspace and full space—to the identity matrix as shown in figure 10. The problem is less trivial for larger  $\Delta t$ . From there, we decided to run various  $\Delta t$  for up to 4.5 seconds.

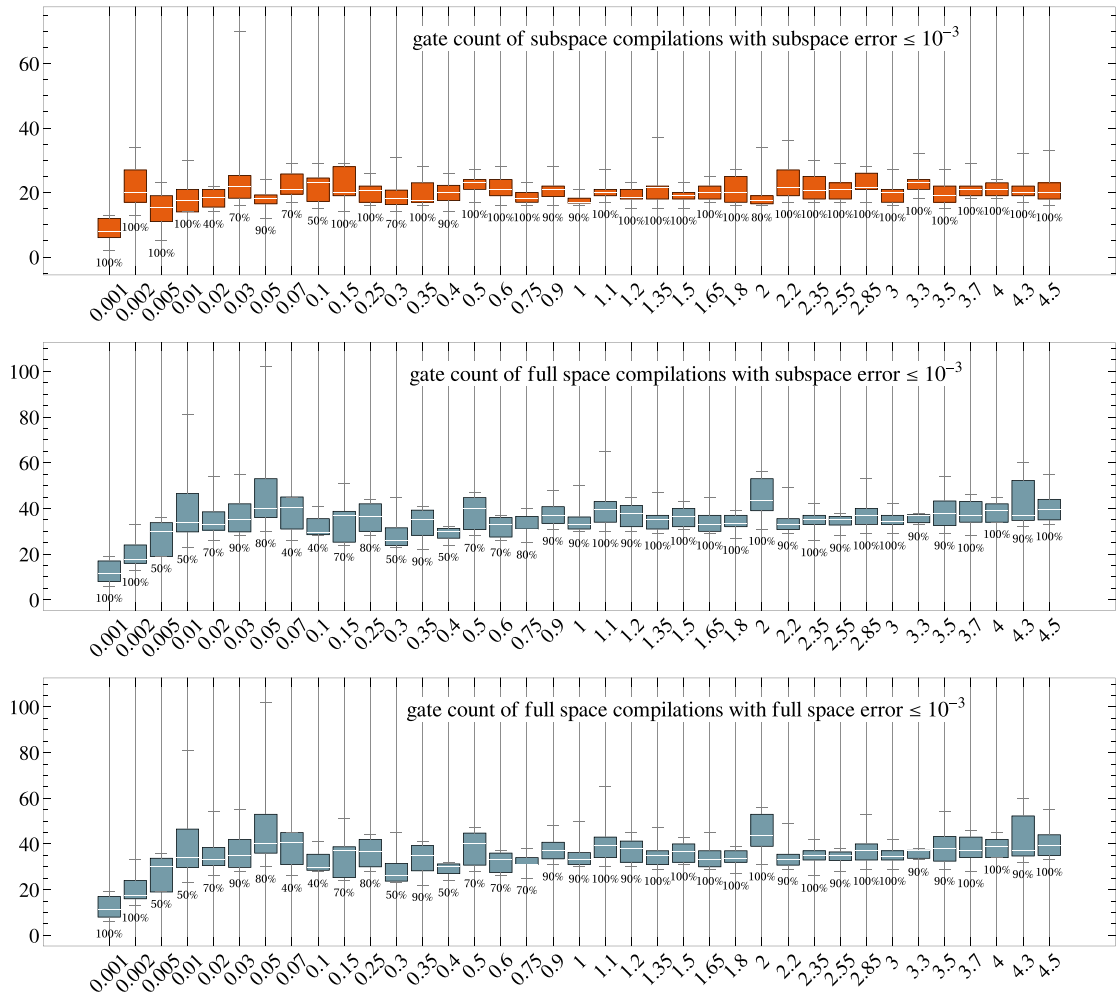
Plotting the best gate count for each cases as of figure 8, but against the distance to the identity matrix, results in figures 11 and 12 for the subspace and full space errors, respectively. These results show that the non-triviality does not correlate with the gate count in the Trotterisation but rather with the  $\Delta t$ .




In synthesising circuits of the time-evolution operators, we run ten trials of circuit synthesis for each  $\Delta t$ , for the subspace and the full space. The computational basis set of our Hydrogen problem is  $\{|i\rangle\}$ , for  $i \in \{0, \dots, 15\}$ . We use the least-significant bit notations, *e.g.*,  $|5\rangle = |0101\rangle$ . Thus, the subspace  $\mathcal{S}$ , the subspace that preserves the number of particles as two is spanned by

$$\text{span}(\mathcal{S}) = \{|0011\rangle, |0101\rangle, |0110\rangle, |1001\rangle, |1010\rangle, |1100\rangle\} \equiv \{|3\rangle, |5\rangle, |6\rangle, |9\rangle, |10\rangle, |12\rangle\}.$$

With that, we obtain the statistic shown in figures 13–16 from our compilations.




**Figure 13.** Gate counts of successful results with two different successful criteria: subspace error  $\leq 10^{-3}$  and full space error  $\leq 10^{-3}$ , for the subspace and full space compilation. The numbers on x-axis denote time durations  $\Delta t$  and the labels below whiskers signifying the success rate of ten trials.

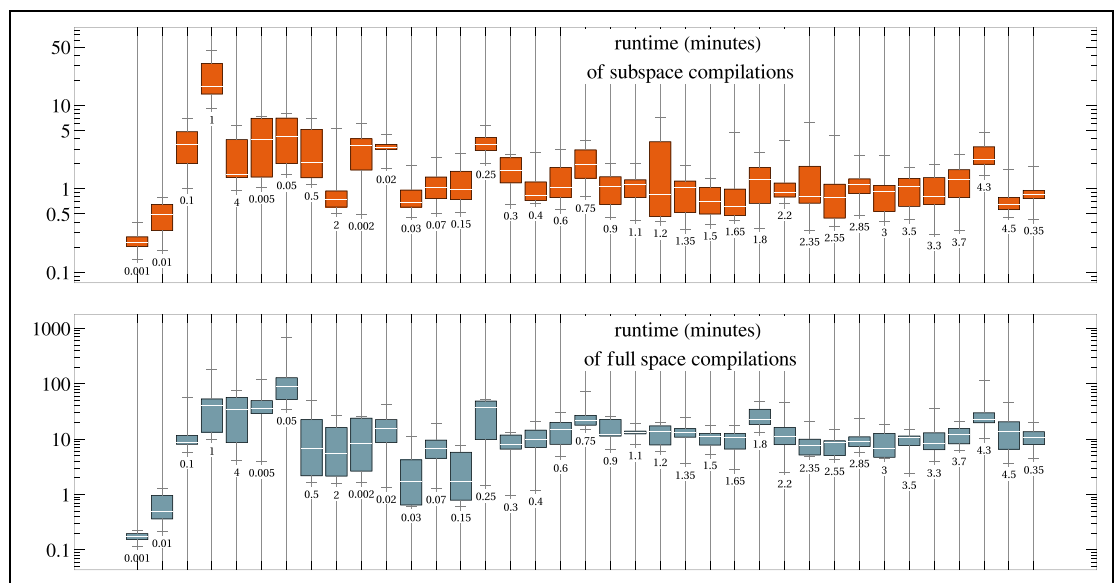


**Figure 14.** Subspace errors. The red dashed lines indicate the goal error:  $10^{-3}$ . Both subspace and full space compilation show similar performance.





**Figure 15.** Full space errors. The red dashed lines indicate the goal error:  $10^{-3}$ . The subspace compilation is obviously perform poorly in this measure since it only takes into account the subspace.



**Figure 16.** The total recompilation runtime for the subspace and full space compilation.

From figure 13, the subspace compilations show lower gate count with more successful rate compared to the full space compilations. The successful results of full space compilations by criteria subspace and full space distance are almost identical with only difference in result of  $\Delta t = 0.75$ .

From figure 14, we see both methods perform relatively similar. However, it is interesting that for some small  $\Delta t$ , the optimiser shows some struggle, in contrast to Trotterisation that perform consistently well for small  $\Delta t$  — see figure 9.

The total optimisation runtime for all compilations shown in figure 16; as we expected, subspace compilations converge in shorter times than the full space compilations.

## ORCID iDs

Cica Gustiani <https://orcid.org/0000-0003-0558-4685>

Richard Meister <https://orcid.org/0000-0002-1998-7867>

Simon C Benjamin <https://orcid.org/0000-0002-7766-5348>

## References

- [1] Preskill J 2018 Quantum computing in the NISQ era and beyond *Quantum* **2** 79
- [2] Peruzzo A, McClean J, Shadbolt P, Yung M-H, Zhou X-Q, Love P J, Aspuru-Guzik A and O'Brien J L 2014 A variational eigenvalue solver on a photonic quantum processor *Nat. Commun.* **5** 1–7
- [3] Kempe J, Kitaev A and Regev O 2005 The complexity of the local hamiltonian problem *FSTTCS 2004: Foundations of Software Technology and Theoretical Computer Science* ed L Kamal and M Meena (Springer) pp 372–83
- [4] Grimsley H R, Economou S E, Barnes E and Mayhall N J 2019 An adaptive variational algorithm for exact molecular simulations on a quantum computer *Nat. Commun.* **10** 1–9
- [5] Ollitrault P J, Baiardi A, Reiher M and Tavernelli I 2020 Hardware efficient quantum algorithms for vibrational structure calculations *Chem. Sci.* **11** 6842–55
- [6] Delgado A, Miguel Arrazola J, Jahangiri S, Niu Z, Izaac J, Roberts C and Killoran N 2021 Variational quantum algorithm for molecular geometry optimization *Phys. Rev. A* **104** 052402
- [7] Gibbs J, Gili K, Holmes Z, Commeau B, Arrasmith A, Cincio L, Coles P J and Sornborger A 2021 Long-time simulations with high fidelity on quantum hardware (arXiv:2102.04313)
- [8] Greene-Diniz G and Ramo D M noz 2021 Generalized unitary coupled cluster excitations for multireference molecular states optimized by the variational quantum eigensolver *Int. J. Quantum Chem.* **121** e26352
- [9] Metcalf M, Bauman N P, Kowalski K and De Jong W A 2020 Resource-efficient chemistry on quantum computers with the variational quantum eigensolver and the double unitary coupled-cluster approach *J. Chem. Theory Comput.* **16** 6165–75
- [10] Hon Sang Chan H, Fitzpatrick N, Segarra-Martí J, Bearpark M J and Tew D P 2021 Molecular excited state calculations with adaptive wavefunctions on a quantum eigensolver emulation: reducing circuit depth and separating spin states *Phys. Chem. Chem. Phys.* **23** 26438–50
- [11] McArdle S, Endo S, Aspuru-Guzik A, Benjamin S C and Yuan X 2020 Quantum computational chemistry *Rev. Mod. Phys.* **92** 015003
- [12] Yifan Li, Jiaqi H, Zhang X-M, Song Z and Yung M-H 2019 Variational quantum simulation for quantum chemistry *Adv. Theory Simulations* **2** 1800182
- [13] Xiaosi X, Sun J, Endo S, Ying Li, Benjamin S C and Yuan X 2021 Variational algorithms for linear algebra *Sci. Bull.* **66** 2181–8
- [14] Endo S, Sun J, Ying Li, Benjamin S C and Yuan X 2020 Variational quantum simulation of general processes *Phys. Rev. Lett.* **125** 010501
- [15] Patil H, Wang Y and Krstić P S 2022 Variational quantum linear solver with a dynamic ansatz *Phys. Rev. A* **105** 012423
- [16] Liu H-L, Wu Y-S, Wan L-C, Pan S-J, Qin S-J, Gao F and Wen Q-Y 2021 Variational quantum algorithm for the poisson equation *Phys. Rev. A* **104** 022418
- [17] Khatri S, LaRose R, Poremba A, Cincio L, Sornborger A T and Coles P J 2019 Quantum-assisted quantum compiling *Quantum* **3** 140
- [18] Jones T and Benjamin S C 2022 Robust quantum compilation and circuit optimisation via energy minimisation *Quantum* **6** 628
- [19] Gokhale P, Ding Y, Propson T, Winkler C, Leung N, Shi Y, Schuster D I, Hoffmann H and Chong F T 2019 Partial compilation of variational algorithms for noisy intermediate-scale quantum machines *Proc. 52nd Annual IEEE/ACM Int. Symp. on Microarchitecture* pp 266–78
- [20] Caro M C, Hsin-Yuan Huang M C, Sharma K, Sornborger A, Cincio L and Coles P J 2021 Generalization in quantum machine learning from few training data (arXiv:2111.05292)
- [21] Cerezo M et al 2021 Variational quantum algorithms *Nat. Rev. Phys.* **3** 625–44
- [22] Tilly J et al 2022 The variational quantum eigensolver: a review of methods and best practices *Phys. Rep.* **986** 1–128
- [23] Fedorov D A, Peng B, Govind N and Alexeev Y 2022 VQE method: a short survey and recent developments *Mater. Theory* **6** 1–21
- [24] McClean J R, Boixo S, Smelyanskiy V N, Babbush R and Neven H 2018 Barren plateaus in quantum neural network training landscapes *Nat. Commun.* **9** 1–6
- [25] Cade C, Mineh L, Montanaro A and Stanisic S 2020 Strategies for solving the Fermi-Hubbard model on near-term quantum computers *Phys. Rev. B* **102** 235122
- [26] Anselme Martin B, Simon P and Rančić M J 2022 Simulating strongly interacting Hubbard chains with the variational hamiltonian ansatz on a quantum computer *Phys. Rev. Res.* **4** 023190
- [27] Kandala A, Mezzacapo A, Temme K, Takita M, Brink M, Chow J M and Gambetta J M 2017 Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets *Nature* **549** 242–6
- [28] Benedetti M, Fiorentini M and Lubasch M 2021 Hardware-efficient variational quantum algorithms for time evolution *Phys. Rev. Res.* **3** 033083
- [29] Rattew A G, Shaohan H, Pistoia M, Chen R and Wood S 2019 A domain-agnostic, noise-resistant, hardware-efficient evolutionary variational quantum eigensolver (arXiv:1910.09694)
- [30] Ho Lun Tang V O S, Barron G S, Grimsley H R, Mayhall N J, Barnes E and Economou S E 2021 Qubit-adapt-VQE: an adaptive algorithm for constructing hardware-efficient ansätze on a quantum processor *PRX Quantum* **2** 020310
- [31] Sim S, Romero J, Gonthier Jôme F and Kunitsa A A 2021 Adaptive pruning-based optimization of parameterized quantum circuits *Quantum Sci. Technol.* **6** 025019
- [32] Chivilikhin D, Samarin A, Ulyantsev V, Iorsh I, Oganov A R and Kyriienko O 2020 MoG-VQE: Multiobjective genetic variational quantum eigensolver (arXiv:2007.04424)
- [33] Meister R, Gustiani C and Benjamin S C 2023 Exploring ab initio machine synthesis of quantum circuits *New J. Phys.* (<https://doi.org/10.1088/1367-2630/ace077>)
- [34] Szabo A and Ostlund N S 2012 *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory* (Courier Corporation)
- [35] Wigner E and Jordan P 1928 Über das Paulische Äquivalenzverbot *Z. Phys.* **47** 631
- [36] Seeley J T, Richard M J and Love P J 2012 The Bravyi-Kitaev transformation for quantum computation of electronic structure *J. Chem. Phys.* **137** 224109
- [37] McClean J R et al 2017 OpenFermion: the electronic structure package for quantum computers (arXiv:1710.07629)
- [38] Sun Q et al 2018 PySCF: the python-based simulations of chemistry framework *Wiley Interdiscip. Rev.-Comput. Mol. Sci.* **8** e1340
- [39] Pople J A 1999 Nobel lecture: quantum chemical models *Rev. Mod. Phys.* **71** 1267
- [40] Hatano N and Suzuki M 2005 Quantum annealing and related optimization methods *Quantum Annealing and Related Optimization Methods* vol XIV, ed A Das and B K Chakrabarti (Springer) p 378

- [41] Sharma K, Khatri S, Cerezo M and Coles P J 2020 Noise resilience of variational quantum compiling *New J. Phys.* **22** 043006
- [42] McArdle S, Jones T, Endo S, Ying Li, Benjamin S C and Yuan X 2019 Variational ansatz-based quantum simulation of imaginary time evolution *npj Quantum Inf.* **5** 1–6
- [43] Stokes J, Izaac J, Killoran N and Carleo G 2020 Quantum natural gradient *Quantum* **4** 269
- [44] Jakob Meyer J 2021 Fisher information in noisy intermediate-scale quantum applications *Quantum* **5** 539
- [45] Jones T and Benjamin S 2020 Questlink-mathematica embiggened by a hardware-optimised quantum emulator *Quantum Sci. Technol.* **5** 034012
- [46] Jones T, Brown A, Bush I and Benjamin S C 2019 Quest and high performance simulation of quantum computers *Sci. Rep.* **9** 1–11
- [47] Meister R 2022 pyQuEST (available at: <https://github.com/rmeister/pyQuEST>) (Accessed August 2022)
- [48] Claudino D, Wright J, McCaskey A J and Humble T S 2020 Benchmarking adaptive variational quantum eigensolvers *Front. Chem.* **8** 1152
- [49] Tkachenko N V, Sud J, Zhang Y, Tretiak S, Anisimov P M, Arrasmith A T, Coles P J, Cincio L and Dub P A 2021 Correlation-informed permutation of qubits for reducing ansatz depth in the variational quantum eigensolver *PRX Quantum* **2** 020337
- [50] Gomes N, Mukherjee A, Zhang F, Iadecola T, Wang C-Z, Kai-Ming H, Orth P P and Yao Y-X 2021 Adaptive variational quantum imaginary time evolution approach for ground state preparation *Adv. Quantum Technol.* **4** 2100114
- [51] Bravyi S, Gambetta J M, Mezzacapo A, and Temme K 2017 Tapering off qubits to simulate fermionic hamiltonians (arXiv:1701.08213)
- [52] Jaeger G 2007 *Quantum Information: An Overview* 1st edn (Springer) pp 94–95
- [53] Barenco A, Bennett C H, Cleve R, DiVincenzo D P, Margolus N, Shor P, Sleator T, Smolin J A and Weinfurter H 1995 Elementary gates for quantum computation *Phys. Rev. A* **52** 3457
- [54] Fleck A Jr, Morris J R and Feit M D 1976 Time-dependent propagation of high energy laser beams through the atmosphere *Appl. Phys.* **10** 129–60
- [55] Kassal I, Jordan S P, Love P J, Mohseni M and Aspuru-Guzik A 2008 Polynomial-time quantum algorithm for the simulation of chemical dynamics *Proc. Natl Acad. Sci.* **105** 18681–6
- [56] Chan H H S, Meister R, Jones T, Tew D P and Benjamin S C 2023 Grid-based methods for chemistry modelling on a quantum computer *Sci. Adv.* **9**