

ARTICLE OPEN



Variational quantum algorithm with information sharing

Chris N. Self¹✉, Kiran E. Khosla¹, Alistair W. R. Smith¹, Frédéric Sauvage¹, Peter D. Haynes^{1,2}, Johannes Knolle^{1,3,4}, Florian Mintert¹ and M. S. Kim¹

We introduce an optimisation method for variational quantum algorithms and experimentally demonstrate a 100-fold improvement in efficiency compared to naive implementations. The effectiveness of our approach is shown by obtaining multi-dimensional energy surfaces for small molecules and a spin model. Our method solves related variational problems in parallel by exploiting the global nature of Bayesian optimisation and sharing information between different optimisers. Parallelisation makes our method ideally suited to the next generation of variational problems with many physical degrees of freedom. This addresses a key challenge in scaling-up quantum algorithms towards demonstrating quantum advantage for problems of real-world interest.

npj Quantum Information (2021)7:116; <https://doi.org/10.1038/s41534-021-00452-9>

INTRODUCTION

Rapid developments in quantum computing hardware^{1–3} have led to an explosion of interest in near-term applications^{4–8}. Though current devices are remarkable feats of engineering, their current coherence times and gate fidelities exclude running general quantum algorithms such as Shor's factorisation, Grover search or quantum phase estimation. Nevertheless, it is hoped that variational quantum algorithms (VQA's) will be able to demonstrate a quantum advantage on noisy intermediate scale quantum (NISQ) devices⁹.

Variational algorithms use low depth quantum circuits as a subroutine in a larger classical optimisation and have been applied broadly, including to binary optimisation problems^{10–12}, training machine learning models^{13–15}, and obtaining energy spectra^{16–18}. While low depth circuits lessen the effect of errors, error rates are still a challenge for current practical implementations. Addressing these errors has been a focus within the literature, with many sophisticated error mitigation approaches being developed^{19–21}. However, errors will not be the only limiting factor for VQA's. Other obstacles include ansatz construction^{22,23}, optimisation challenges^{24,25} and integrated hardware design²⁶.

VQA's are very demanding of quantum hardware, requiring large numbers of sequential calls to quantum devices. In many cases demands on device throughput are further exacerbated by a need to solve multiple different but related optimisation problems (e.g. molecules with different nuclear separations¹⁶, or edge-weighted graphs for different weights¹²), which to date have mostly been treated independently. Few notable exceptions have solved related problems iteratively²⁷. Parallelising related problems offers a way to maximize the utility of each circuit. This is particularly relevant to cloud-based interfaces—rapidly becoming the industry standard—however, even with dedicated device access, the low-throughput of quantum hardware remains a limiting factor, worsening for larger and more complex problems.

In this paper, we introduce a parallel optimisation scheme for VQA problems where the cost function is parameterised by some physical parameter(s). A collection of problems, corresponding to different values of physical parameters, are optimised in parallel by an array of Bayesian optimisers sharing quantum measurement results between them. Using our Bayesian optimisation with

information sharing (BOIS) approach, we demonstrate a significant reduction in the number of circuits required to obtain a good solution at all parameter points. BOIS allows us to efficiently find potential energy curves and surfaces for small molecules and a quantum spin model on IBM Quantum devices²⁸.

RESULTS

Parallelising VQE

As an example VQA, we consider the variational quantum eigensolver (VQE)²⁹, used to find the ground states of qubit Hamiltonians H . VQE uses a quantum ansatz circuit, parameterised by a set of angles θ , to prepare a state $|\psi(\theta)\rangle$, and estimate the expectation value $\langle H \rangle_{\theta} \equiv \langle \psi(\theta) | H | \psi(\theta) \rangle$. Classical optimisation is used to minimize the cost function $\langle H \rangle_{\theta}$ with respect to θ , which—for a sufficiently expressive ansatz (see “Ansatz design” in the “Methods” section)—corresponds to the ground state.

Qubit Hamiltonians can always be decomposed into a weighted sum of Pauli strings $P_i \in \{I, X, Y, Z\}^{\otimes n}$ (for n qubits) with the weights c_i depending on some physical parameter(s) \mathbf{x} ,

$$H(\mathbf{x}) = \sum_i c_i(\mathbf{x}) P_i. \quad (1)$$

These physical parameters are, for example, nuclear coordinates of a molecule, length scale of long-range interactions, distance of reactants above a catalyst, etc.

From the structure of Eq. (1), it is clear that optimising for each \mathbf{x} independently does not exploit the relationships between problems at different physical parameters. In particular, measuring the required Pauli expectation values $\langle P_i \rangle_{\theta}$ at a single θ point can be used to estimate $\langle H(\mathbf{x}) \rangle_{\theta}$ for all \mathbf{x} , as it is simply a weighted sum of Pauli expectation values with weights $c_i(\mathbf{x})$

Bayesian optimisation and BOIS

In order to exploit the relationships between VQE problems, we use an optimisation method that collates function evaluations across the entire θ parameter space. Bayesian optimisation (BO) uses a surrogate model of the global cost function to guide the optimization³⁰. The surrogate model at unevaluated θ points is inferred, by means of Bayes rules, based on all previous cost

¹Blackett Laboratory, Imperial College London, London, UK. ²Department of Materials, Imperial College London, London, UK. ³Department of Physics TQM, Technische Universität München, Garching, Germany. ⁴Munich Center for Quantum Science and Technology (MCQST), Munich, Germany. ✉email: cns08@ic.ac.uk

functions evaluations. Every optimization round, the next θ point is decided by maximizing a utility function over the surrogate model (see “Practical aspects of Bayesian Optimisation” in the “Methods” section). BO has been applied in quantum control^{31–34}, as well as in experimental variational algorithms^{35,36}, however its use in VQE and the exploitation of its global surrogate model remains largely unexplored.

BOIS employs an array of Bayesian optimisers running in parallel. Each optimiser, $\mathcal{B}(\mathbf{x}_a)$, attempts to solve a separate but related VQE problem, by minimising $\langle H(\mathbf{x}_a) \rangle_{\theta}$ over θ at physical coordinate \mathbf{x}_a , see Fig. 1. At each iteration, optimiser $\mathcal{B}(\mathbf{x}_a)$ requests a cost function evaluation at the Bayes optimal point to evaluate next θ_a (see “Practical aspects of Bayesian optimisation” in the “Methods” section). As described in the section “Parallelising VQE”, the Pauli string measurement results required to evaluate $\langle \psi(\theta_a) | H(\mathbf{x}_a) | \psi(\theta_a) \rangle$ also allows us to evaluate $\langle \psi(\theta_a) | H(\mathbf{x}_\beta) | \psi(\theta_a) \rangle$, which can be passed to the optimiser $\mathcal{B}(\mathbf{x}_\beta)$. We refer to these cross-evaluations as information sharing from $\mathcal{B}(\mathbf{x}_a)$ to $\mathcal{B}(\mathbf{x}_\beta)$. This requires each optimiser to

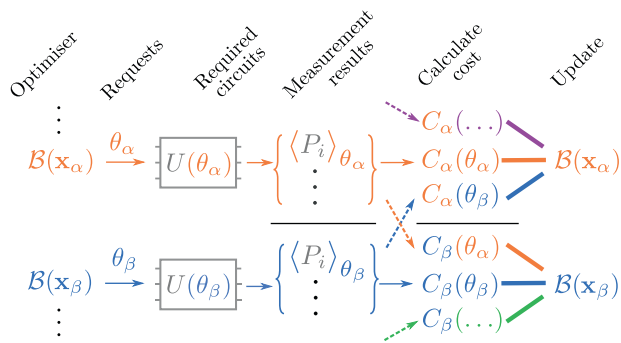


Fig. 1 Information sharing framework. Parallel Bayesian optimisation for physically parameterised VQE tasks. Separate BOs $\mathcal{B}(\mathbf{x}_a)$ optimise for different cost functions C_a , corresponding to different values of the physical parameters \mathbf{x}_a , using the same parameterised ansatz circuit $U(\theta)$. Every iteration, each $\mathcal{B}(\mathbf{x}_a)$ requests a new variational parameter point θ_a , at which the set of Pauli strings $\{P_i\}$ are measured. These expectation values are used to compute any C_β cost functions (dashed lines) at θ_a for all a, β . Each BO can then be updated using the measurement results obtained for several $\theta_a, \theta_\beta, \dots$ parameter points each iteration (bold arrows), dramatically speeding up convergence at all \mathbf{x}_a .

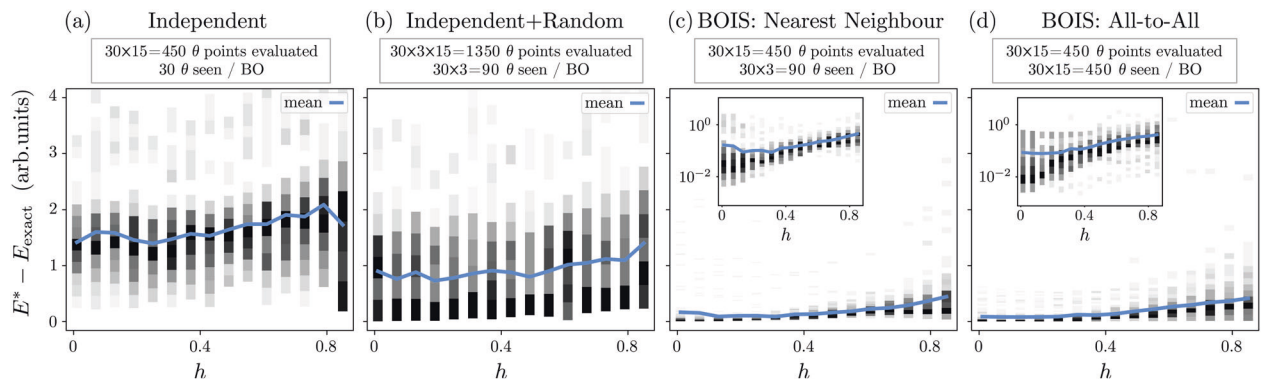


Fig. 2 Numerical study of the effectiveness of information sharing. Comparison between (a–b) independent vs (c–d) information shared VQE optimisation strategies for the quantum spin model, Eq. (2) with field $h = h_x = h_z$. The sample distribution of 100 repetitions of the same optimisation are shown as density plots for each h_a , with darker shades corresponding to more observations, and the solid curve showing the mean. Optimisation is run for thirty iterations and we consider fifteen values of h_a . Data boxes show the total number of function evaluations per repetition and the number of evaluations seen by each optimiser $\mathcal{B}(h_a)$, in addition to the numbers quoted each BO receives ten initialisation data points that are also either independent or shared (see “Practical aspects of Bayesian optimisation” in the “Methods” section). **a** Independent BOs at each h_a . **b** As before, but each optimiser requests two additional energy evaluations (at random θ parameter points) at each iteration. **c** BOIS nearest-neighbour information sharing, i.e. BOs at neighbouring h -field points (e.g. h_a and $h_{a\pm 1}$) share device measurement results. **d** BOIS all-to-all information sharing, i.e., every BO sees all device measurement results. Insets of (c) and (d) show the same data on a log-scale.

use the same θ -parameterised ansatz circuit and that $H(\mathbf{x}_a)$ decomposes into the same set of Pauli strings $\{P_i\}$ for all a (this can be padded as necessary).

The intrinsic locality of gradient-based approaches limits the amount of information that can be shared between different optimization runs. While we focused on BO, other surrogate model-based approaches³⁷ could be equivalently lifted to benefit from this information sharing scheme.

Testing the effectiveness of information sharing

Firstly we demonstrate the effectiveness of information sharing by using BOIS for a VQE task (VQE + BOIS) applied to a quantum spin chain with Hamiltonian

$$H(h_x, h_z) = \sum_{\langle i,j \rangle} Z_i Z_j - \sum_i (h_x X_i + h_z Z_i). \quad (2)$$

The (dimensionless) transverse h_x and longitudinal h_z fields are classical coordinates analogous to nuclear separation in molecules. For $h_x > 0$ and $h_z = 0$, Eq. (2) can be diagonalized by a Jordan-Wigner transformation, which transforms the spins into non-interacting fermions. For $h_x, h_z > 0$ it becomes non-integrable, with approximate methods eventually breaking down around a phase transition occurring along a critical line³⁸.

We consider the case where $h_x = h_z = h \in [0, 0.9]$, which we discretise into 15 values h_a . This range of h approaches but does not cross the critical line. The VQE is then run using our BOIS optimisation approach for a system of four spins, with open boundary conditions. We use an ansatz circuit systematically generated—using an original approach (see “Practical aspects of Bayesian optimisation” in the “Methods” section)—for this Hamiltonian. Noiseless simulations are carried out by contracting tensor network representations of the ansatz circuit and Hamiltonian using the Quimb python package³⁹.

The performance of four different information-sharing strategies is shown in Fig. 2. Firstly, as a baseline, we consider independent BO with no information sharing Fig. 2a, followed by independent BOs with two extra (random) function evaluations per iteration Fig. 2b. This second case is chosen such that the number of evaluations seen by each $\mathcal{B}(h_a)$ is the same as nearest-neighbour information sharing (each h_a shares only with $h_{a\pm 1}$) Fig. 2c. Finally, we show all-to-all information sharing (each h_a shares with all $h_\beta, \beta \neq a$) Fig. 2d, which requires the same number of

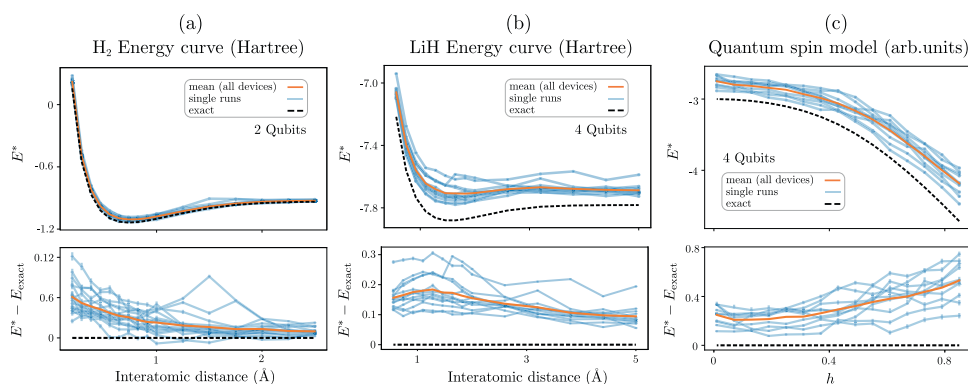


Fig. 3 Experimental application of VQE + BOIS running on IBMQ devices. Final energy estimates E^* (top) and final errors ($E^* - E_{\text{exact}}$) (bottom) are plotted for three different problems. Repeated runs (blue curves), executed on a set of different devices²⁸, show the distribution around the mean (orange curves). **a** H₂ dimer showing eighteen separate VQE+BOIS (all-to-all) runs. Each ran for ten iterations with thirty (shared—see “Practicals aspects of Bayesian optimisation” in the “Methods” section) initial points. **b** LiH dimer showing thirteen separate VQE +BOIS (all-to-all) runs. Each ran for thirty iterations with thirty (shared) initial points. **c** Quantum spin model, Eq. (1) with $h_x = h_z = h$, (considered in “Testing the effectiveness of information sharing”) showing ten separate VQE + BOIS (nearest-neighbour) runs. Each ran for fifty iterations with ten (shared) initial points. Error bars indicate standard error on final energy measurements.

function evaluations as nearest-neighbour sharing but employs more cross-evaluations. Each strategy is run for thirty iterations, after which we compare the final VQE energy estimate (E^*) to the exact ground state energy (E_{exact}). Optimisation of the full energy curve is repeated 100 times for each strategy. We present the data as a shaded histogram of $E^* - E_{\text{exact}}$ at each h_{or} as well as the sample mean. Each optimisation run is initialised with ten cost function evaluations, these are either shared for the BOIS strategies Fig. 2c–d or independent Fig. 2a–b, as discussed in “Practicals aspects of Bayesian optimisation” in the “Methods” section.

We see clear improvement of sharing strategies over independent optimisations. After thirty iterations none of the independent BO Fig. 2a cases have converged close to the exact ground state energy. Adding extra energy evaluations (at random θ) to the independent optimisations improves their performance Fig. 2b, highlighting the capacity of BO to leverage the full dataset of evaluations it can access. Despite the improvement, a significant portion of runs end far from the ground state, with non-zero mean error and large variance. Nearest-neighbour sharing Fig. 2c hugely improves on Fig. 2b despite the individual BO’s of the two strategies receiving the same number of function evaluations per iteration. Here, the extra evaluations have greater utility compared to Fig. 2b as they come from BO’s solving similar (i.e. for close values of h) optimisation problems. Finally, All-to-All sharing shows moderate improvement over nearest-neighbour sharing, Fig. 2d in spite of having received five times more energy evaluations. This indicates that most of the advantage of the information sharing is coming from BO’s that are close together in h . An additional advantage of nearest-neighbour sharing is it reduces the computational overhead of updating the surrogate model by only including the most relevant shared data. As expected, we find the hardest optimisation problems, where our sharing strategies still do not always converge, to be large values of h closest to the critical line.

Experimental demonstration on IBMQ devices

In the following, we demonstrate the success of the BOIS strategy on problems run on IBMQ devices. By studying a quantum spin model (used in “Testing the effectiveness of information sharing”) as well as H₂, LiH and a linear chain of H₃, we highlight the capability of our approach to complete VQE tasks in practical timescales. Optimisation tasks were carried out either on Paris, Toronto, Athens, Manhattan, Valencia, or Santiago IBMQ quantum processors²⁸. We employ Qiskit’s⁴⁰ built in ‘Complete

Measurement Filter’ with no further error mitigation. During optimisation 1024 measurement shots are used, increasing to 8192 shots for the final reported VQE energies.

Figure 3 shows the BOIS methods applied to finding ground states of H₂ (all-to-all sharing), LiH (all-to-all sharing) and the spin model (nearest-neighbour sharing), each of which only has one physical parameter. These are two, four and four qubit problems respectively, with six, thirteen, and ten optimisation parameters respectively (see “Practicals aspects of Bayesian optimisation” in the “Methods” section). In obtaining the Hamiltonians for H₂ and LiH, we have removed the spin \mathbb{Z}_2 symmetries, and in the case of LiH, have also frozen non-participating orbitals (see “Qubit Hamiltonians for small molecules” in “Methods” section). H₂ can be reduced to a single qubit⁴¹, however, for comparison with previous work, the two-qubit Hamiltonian is used. These optimisations converge after a remarkable ~10–50 iterations, far below what could be reasonably expected for any stochastic gradient descent, which would require many hundreds of iterations. Such a drastic improvement unequivocally demonstrates the advantage of data sharing between optimizers, showing just how much information can be extracted from each set of quantum measurements.

We now consider a linear chain of H₃, parameterized by two relative inter-atomic distances $\mathbf{x} = (x_1, x_2)$, which we discretize into a grid to find the 2D energy surface. Figure 4 shows the two-dimensional potential energy surface found using VQE + BOIS with nearest-neighbour sharing. For this two-dimensional problem nearest-neighbour sharing means (x_1^a, x_2^b) shares with (x_1^{a+1}, x_2^b) , (x_1^{a-1}, x_2^b) , (x_1^a, x_2^{b+1}) and (x_1^a, x_2^{b-1}) . Solving for the eight-by-eight grid in parallel means 64 θ points per iteration, however, we construct the entire surface with just ten iterations, for a total of 650 individual θ evaluations (including initialisation). In contrast, we expect gradient descent (in this ten-dimensional optimisation space) would require a similar number of θ evaluations just for a single physical coordinate¹⁶.

DISCUSSION

Our experimental results, Fig. 3, show a discrepancy between experimental and theoretical energy curves/surfaces. The most significant sources of VQE errors are (i) optimisation errors, and (ii) device errors. Ansatz expressibility errors, the other limiting factor for VQE, can be ruled out by our custom circuit design (see “Ansatz design” in the “Methods” section). Known difficulties in the optimisation landscapes mean single runs of the optimisation will

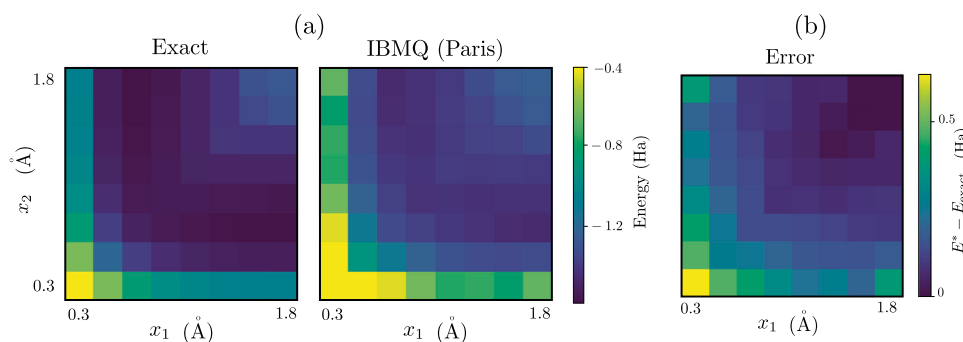


Fig. 4 Experimental application to a multi-dimensional problem. Ground state energies (in Hartree) for a linear chain of H_3 , over two-dimensional parameter space (x_1, x_2) where x_1 the distance between the first and second H atoms and x_2 the distance between the second and third. **a** the exact ground state energy surface for an 8×8 grid of x_1 and x_2 values, compared to experimental results from running VQE + BOIS on from IBM's Paris chip. The optimisations converged in ten iterations to compute the entire surface. **b** the error in the VQE + BOIS results, showing systematic offsets due to gate errors.

sometimes fail²⁴. This is exemplified in Fig. 2 for large h where the ground state is more entangled. The same qualitative behaviour is observed in the experimental data Fig. 3c and is made worse by the randomness from using a noisy quantum device to estimate the cost function.

The systematic offset in the experimental data, Fig. 3, is typical of depolarising errors in quantum gates, becoming worse for more qubits¹⁹. Variations between runs are expected as different devices have different error characteristics and additionally these change over time. In general, the variance is further compounded by shot noise (i.e. finite numbers of measurement on device), but this is well below the variance of device noise across different days and devices.

Why would we expect this information sharing to be helpful? Considering a molecule as an example, a small change in the nuclear coordinates can be understood as a perturbation of the electronic Hamiltonian. The perturbed and unperturbed ground states will have some overlap, so their optimal θ 's will be found in a similar region of parameter space. We can then see that a useful θ_a selected by optimiser $\mathcal{B}(\mathbf{x}_a)$ is likely to be valuable to $\mathcal{B}(\mathbf{x}_a + \delta\mathbf{x})$. In this way, information sharing allows each BO to exploit the promising regions of θ -space discovered by other optimisers, making the evaluations of adjacent optimisers more valuable than random θ evaluations.

Information sharing would be particularly helpful when the physical parameter tunes some non-trivial interaction that continuously increases the entanglement in the ground state, for example approaching a phase transition as in "Testing the effectiveness of information sharing". In this case we would expect the optimisation to converge faster in the simple (less entangled ground state) limit and this information to flow to harder (more entangled) problems.

Here we introduced BOIS, an efficient optimisation scheme for related VQA's. Our approach is specifically designed to make maximum use of limited numbers of quantum measurements in order to combat the limited throughput of NISQ quantum devices. We demonstrated the efficiency of our approach by experimentally solving for energy curves in tens, instead of hundreds or thousands of iterations. The speedup provided by our method has enabled both 1D energy curves and the 2D energy surface of a molecular trimer to be found on a practical timescale, even without dedicated device access. BOIS can be readily applied to other parameterised VQA's and, combined with our circuit construction, makes benchmarking VQA algorithms on real devices significantly more practical.

METHODS

Qubit Hamiltonians for small molecules

To find molecular ground states we first reformulate the problem from electrons to qubits. This is discussed at length in the literature, e.g., in the reviews^{7,8}. Here we outline the general procedure and highlight features relevant to the above simulations.

The electronic states are expanded in hydrogen-like orbitals centred on each atom, specifying a fermionic operator and a spatial wave function for each spin orbital. The electronic spatial dependence can be integrated out under the Born-Oppenheimer approximation. This re-expresses the kinetic and potential energy terms as non-linear coupling rates between the fermionic operators, where the couplings rates are nuclear coordinates dependent. The Hilbert space can be further reduced by removing or freezing (i.e. permanently occupied) spin orbitals. Finally, the anti-commuting fermionic operators are transformed into commuting Pauli operators using the Jordan-Wigner, Bravi-Kitaev, or some similar map⁷.

In this work we use Openfermion⁴² (for H_3) and Qiskit⁴⁰ (for H_2 and LiH) to compute the one and two body integrals, using the STO-3G approximation. For H_2 and H_3 , only the 1s orbitals are used (none of which are frozen), and the systems have a total of two and three electrons respectively. For LiH, the Li:1s, orbital is doubly occupied with electrons and frozen, while the four Li:2p_y, Li:2p_z spin orbitals are frozen with zero occupation. Thus only the four Li:2s,2p_x spin orbitals (bond axes aligned with x) and two H:1s spin orbitals constitute the available states for the remaining two electrons. We use the symmetry conserving Bravi-Kitaev transform for H_3 , and the parity transforms for H_2 and LiH to express the fermionic Hamiltonian as a qubit Hamiltonian. For each molecule the \mathbb{Z}_2 symmetries (corresponding to particle and spin symmetries) have been removed^{16,41}, reducing the number of qubits by two.

Ansatz design

As our focus is on efficient optimisation, rather than the challenging and distinct problem of blind ansatz construction, we construct ansatzes that are tailored to the problems we consider. The effectiveness of VQE is highly dependent on the suitability of the ansatz used. An ansatz circuit must be sufficiently expressive to produce its target state, however this must be balanced against other practical concerns including hardware limitations (restricted qubit connectivity and practical limits on circuit depth due to experimental noise) as well as optimisation considerations (the number of optimisation parameters must not be too large).

Our ansatzes, for systems other than H_2 , are constructed classically such that they are guaranteed to be able to produce the target states with high fidelity (in noiseless simulations), while being both depth and parameter number efficient. This allows us to assess the performance of our optimisation scheme without it being limited by the expressive power of the ansatz. In addition, this allows us to bake hardware constraints, such as qubit connectivity, into the construction process. The ansatzes we use for each of the systems we consider are shown in Fig. 5.

Building the ansatz circuit is done in two phases; a growth phase in which the expressivity of the circuit is increased, and a shrinkage phase in which redundant parameters are removed. This approach is inspired by

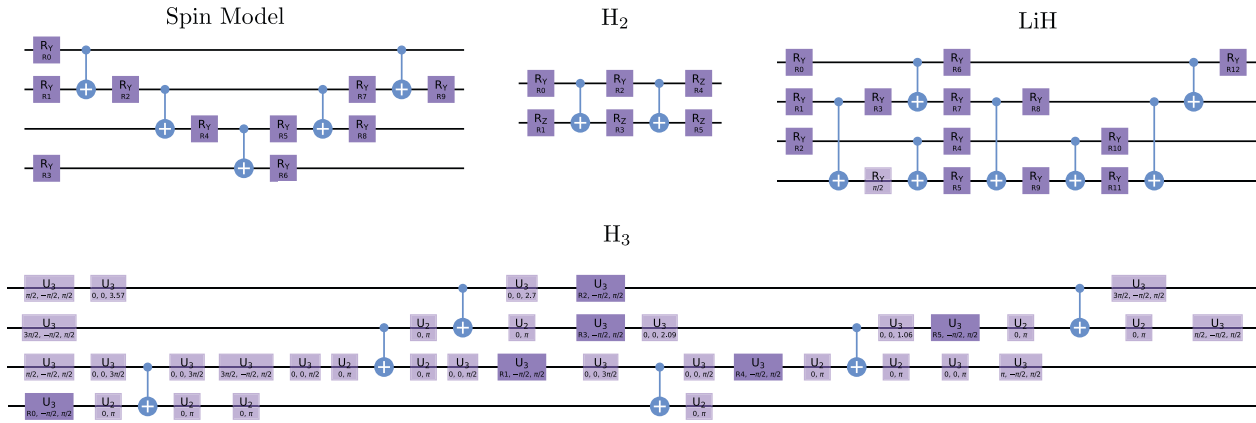


Fig. 5 Parameterised quantum circuits used as ansatz. Ansatz circuits used for each of the physical systems we consider, generated using the procedure described in “Ansatz design”. Gates referenced are from the Qiskit gate set⁴⁰. Parameterised gates are shown in solid colours and fixed angle rotations are indicated with partial transparency.

qubit-ADAPT²³ but uses exhaustive optimisation to decide how to grow the circuit rather than a gradient-based condition. The noiseless simulations used to construct these ansatzes are carried out with highly efficient tensor-network calculations⁴³, performing gradient-descent with the Quimb Python package³⁹. We now describe each phase in detail.

The growth phase begins with a separable circuit consisting only of single-qubit rotations and adds two-qubit entangling blocks in the locations that provide the greatest increase in the fidelity with the target state. Entangling blocks consist of a two-qubit gate, e.g. CNOT, padded on either side with parameterised single-qubit rotations. In general, the single-qubit gates are arbitrary Bloch sphere rotations (e.g. Qiskit’s U_3 gate or equivalent) to allow maximum flexibility. However, in many cases simpler gates can be used (such as real-valued Hamiltonians where R_Y is sufficient). At each step, we determine all locations that it is possible to place a new entangler, compatible with the qubit connectivity. Early in the optimisation this is a weak constraint as we do not initially enforce a mapping of virtual to physical qubits. As the ansatz grows the mapping is fixed by the previously chosen locations and this greatly reduces the number of potential locations in subsequent steps. We test each possible location by optimising the fidelity \mathcal{F} between the target state and the circuit with this block added. For small to moderately sized systems, noiseless tensor-network-based gradient descent can be performed efficiently with automatic differentiation (e.g. using Google’s TensorFlow⁴⁴) to minimise a cost \mathcal{C} . For a target state $|\psi_t\rangle$ and trial state $|\psi\rangle$ we used

$$\mathcal{C} = 1 - \mathcal{F} = 1 - |\langle\psi|\psi_t\rangle|. \quad (3)$$

(Here we use this convention for the fidelity, as opposed to $|\langle\psi|\psi_t\rangle|^2$, as this yielded better performance.) The entangling block that leads to the greatest increase in fidelity is added to the ansatz and the process is repeated. Since our scheme relies on direct optimisation we are able to trial placing entangling blocks at both the beginning and end of the circuit. Placing an entangling block at the beginning of a circuit can produce a more dramatic change than placing one at the end, potentially allowing our method to converge to the target state more quickly than other greedy methods such as qubit-ADAPT. The ansatz that is produced in the growth phase is very efficient in terms of depth. Using gradient descent the growth stage can be performed with relatively few cost/gradient evaluations.

Once the circuit has converged to an acceptable fidelity with the target state (we used $1 - \mathcal{F} < 10^{-6}$) we begin the shrinkage phase. We re-optimize the fidelity of the ansatz with respect to the target state but now adding a regularisation penalty to our cost function which, for gate angles $\{\phi_i\}$, becomes

$$\mathcal{C} = 1 - \mathcal{F} + \eta \sum_i D(\phi_i, 2\pi) \quad (4)$$

Here η is a small regularisation weight parameter and $D(\phi_i, 2\pi)$ is the absolute distance between ϕ_i and its nearest multiple of 2π . This is an example of (periodic) L_1 regularisation and, provided η is small, encourages gate angles to shrink towards a multiple of 2π while maintaining a high fidelity. This allows us to identify unnecessary single-qubit gates (those with angles close to multiples of 2π) which are then removed from the

ansatz. This regularisation process is repeated with increasing η until the point where removing any more single-qubit gates would appreciably degrade the fidelity (checking if this is the case by re-optimising the fidelity using the ansatz with small parameter gates removes).

The ansatz construction process is non-deterministic however for the systems considered we found that typically required numbers of cost/gradient evaluations were (for the growth/shrinkage stages respectively) – H_3 :~4500, 1200 (using general U_3 gates); LiH~1800, 600; quantum spin model~1800, 500.

These ansatzes are produced by optimising the fidelity with a single target state. In our VQE + BOIS simulations we are attempting to find ground state energy surfaces across a physical parameter space. For each VQE + BOIS simulation a target state was chosen as a ground state sitting on the energy surface in question. However, these target states were chosen such that their physical parameter \mathbf{x} was not one that was associated with a BOIS optimiser $\mathbf{x} \notin \mathbf{x}_g$. Although our ansatzes were produced targeting just one state on each energy surface, we found that provided it can represent this target state with a high fidelity (around $1 - \mathcal{F} < 10^{-6}$) then it will typically have similar performance on the other states of the surface (assuming we do not cross any complicated quantum phase transitions). Finally, to further reduce the number of optimisation parameters, the optimal gate parameters for each BOIS physical parameter \mathbf{x}_g are found (again by maximising the fidelity) and any single-qubit gate angles within this optimal parameter set that are found to remain effectively constant across the energy surface are fixed to these constant values in the ansatz. These fixed angle gates are indicated with partial transparency in Fig. 5.

While this ansatz preparation procedure uses the fidelity with a known target state as a cost function, requiring knowledge of the target state, this is not an a priori requirement. With minimal adjustments the same tensor-network-based framework can be adapted to use the expectation of a Hamiltonian as a cost function to be minimised, resulting in a scheme that is qualitatively similar to hardware-efficient ADAPT-VQE²³. An energy-based cost function would be less specific than one based on the fidelity with a target state and so convergence issues may arise if the Hamiltonian possesses many near-degeneracies. The lifting of degeneracies across physical parameter space may also prevent an energy-optimised ansatz from generalising well. Ultimately these classical gradient-based methods are limited to systems of only up to around 10 qubits (we have tested up to this size).

As we aimed to build tailored VQE ansatzes to benchmark BOIS, rather than develop a full scalable adaptive ansatz algorithm, we employed the more reliable fidelity-based approach. However, we believe an adaptive scheme in which an ansatz is grown ansatz one entangling block at a time followed by parameter regularisation has the potential to be run in a more scalable way on-device. Each stage of such a process would be effectively its own VQE problem, potentially allowing information sharing to increase its efficiency, provided that the optimisation is done globally such that information can be shared. By considering all physical parameter points when choosing entangling block placement and performing regularisation, a general ansatz for the whole energy surface may also be obtained more easily than when considering only a single physical parameter point.

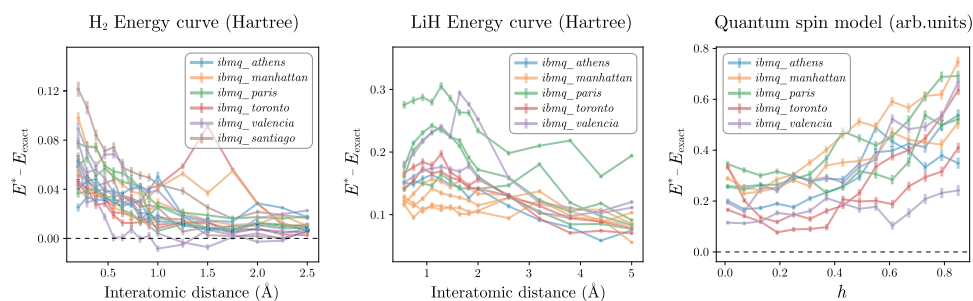


Fig. 6 Experimental data separated by IBM Quantum device. Final errors, $E^* - E_{\text{exactr}}$ in repeated executions of VQE + BOIS ground state energy estimates for H_2 , LiH and the quantum spin model with the IBM Quantum device used indicated by plot colour. We do not observe any clear, systematic differences between the different devices. Error bars indicate standard error on final energy measurements.

Practicals aspects of BO

In our work BO is implemented using the GPyOpt python package⁴⁵. We have used the standard multipurpose Matern-5/2 kernel. Whether more specialised kernels are more suitable for VQA's is an interesting question.

BO begins with a pool of initial data. These are cost functions evaluations made at random in the optimisation parameter space, that are used to initialise the surrogate model. We select the M initial points by Latin hypercube sampling. When considering our information-sharing strategies we additionally share these M initial points amongst the different optimisers $\mathcal{B}(\mathbf{x}_i)$. In contrast, when we compare this to independent strategies with no information sharing each optimiser $\mathcal{B}(\mathbf{x}_i)$ sees a different set of M initial points.

In all of our results, we use a lower-confidence bound (LCB) acquisition function⁴⁶, defined by

$$a_{\text{LCB}}(\theta) = \mu(\theta) - \kappa \sigma(\theta),$$

where μ and σ^2 are the mean and variance functions of the Gaussian process surrogate model. The hyperparameter κ is chosen to decrease linearly from an initial value κ_0 to zero at the final iteration N , so that at iteration $t \in [1, N]$ it has value

$$\kappa_t = \kappa_0 \frac{N - t}{N}.$$

This highly weights exploration early in the optimisation, while prioritising exploitation later on. The value of κ_0 is tuned between one and five depending on the problem. At iteration t the next θ to evaluate is selected by minimising $a_{\text{LCB}}(\theta)$.

BO is conventionally limited to small numbers of parameters, typically $<20-30$ ³⁰. This will become a problem for larger VQE problems, especially when using hardware efficient ansatz. However, decreasing noise rates will help alleviate this issue, since it will allow more complex but parameter efficient ansatz, such as unitary coupled cluster (UCC)⁷, to be employed. Beyond this, global optimisation with large numbers of parameters is an active research topic. Replacing the Gaussian Process with other surrogate models, such as neural networks or random forests, can potentially allow much larger numbers of parameters⁴⁷.

Problems such as over-parameterisation can, in principle, be solved by using physically motivated ansatz circuits, such as UCC for molecules⁴⁸, or schemes that grow an ansatz systematically, for example, ADAPT-VQE^{22,23}. These approaches typically require circuits that are too deep to be practically run on current NISQ devices, however more sophisticated schemes such as ansatz that preserve symmetries⁴⁹ may help to overcome these challenges.

Separating experimental data by device

The experimental data presented in Fig. 3 was collected in October and November 2020 from the IBM quantum devices Athens, Manhattan, Paris, Toronto and Valencia, with additional H_2 data from Santiago²⁸. In Fig. 6 the final VQE errors, $E^* - E_{\text{exactr}}$ is plotted for the H_2 , LiH and quantum spin model tasks, with the IBM device the data was collected from indicated by plot colour.

All of the devices we use have quantum volume 32, with the exception of Valencia which has quantum volume 16. Within our data, we do not see clear systematic differences between the devices when looking across all the tasks. For a given task there can appear to be clustering. For example, in our LiH data, Paris appears to do consistently worse and Manhattan better. However, this is more likely related to calibration differences in the

time window of use or small sample sizes. It is interesting to note, though, that the only non-physical results (ground state energies lower than the true ground state energy) occur for the H_2 task on the lower quantum volume device, Valencia.

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

CODE AVAILABILITY

The code needed to reproduce the findings of this study is available from the corresponding author upon reasonable request.

Received: 5 February 2021; Accepted: 23 June 2021;
Published online: 22 July 2021

REFERENCES

- Jurcevic, P. et al. Demonstration of quantum volume 64 on a superconducting quantum computing system. *Quantum Sci. Technol.* **6**, 025020 (2021).
- Arute, F. et al. Quantum supremacy using a programmable superconducting processor. *Nature* **574**, 505–510 (2019).
- Pino, J. et al. Demonstration of the trapped-ion quantum ccd computer architecture. *Nature* **592**, 209–213 (2021).
- Preskill, J. Quantum computing in the nisq era and beyond. *Quantum* **2**, 79 (2018).
- Benedetti, M., Lloyd, E., Sack, S. & Fiorentini, M. Parameterized quantum circuits as machine learning models. *Quantum Sci. Technol.* **4**, 043001 (2019).
- Smith, A., Kim, M., Pollmann, F. & Knolle, J. Simulating quantum many-body dynamics on a current digital quantum computer. *npj Quantum Inf.* **5**, 1–13 (2019).
- McArdle, S., Endo, S., Aspuru-Guzik, A., Benjamin, S. C. & Yuan, X. Quantum computational chemistry. *Rev. Mod. Phys.* **92**, 015003 (2020).
- Bauer, B., Bravyi, S., Motta, M. & Kin-Lic Chan, G. Quantum algorithms for quantum chemistry and quantum materials science. *Chem. Rev.* **120**, 12685–12717 (2020).
- Cerezo, M. et al. Variational quantum algorithms. Preprint at <https://arxiv.org/abs/2012.09265> (2020).
- Farhi, E., Goldstone, J. & Gutmann, S. A quantum approximate optimization algorithm. Preprint at <https://arxiv.org/abs/1411.4028> (2014).
- Harrigan, M. P. et al. Quantum approximate optimization of non-planar graph problems on a planar superconducting processor. *Nat. Phys.* **17**, 332–336 (2021).
- Zhou, L., Wang, S.-T., Choi, S., Pichler, H. & Lukin, M. D. Quantum approximate optimization algorithm: performance, mechanism, and implementation on near-term devices. *Phys. Rev. X* **10**, 021067 (2020).
- Havlíček, V. et al. Supervised learning with quantum-enhanced feature spaces. *Nature* **567**, 209–212 (2019).
- Grant, E. et al. Hierarchical quantum classifiers. *npj Quantum Inf.* **4**, 1–8 (2018).
- Schuld, M., Sweke, R. & Meyer, J. J. Effect of data encoding on the expressive power of variational quantum-machine-learning models. *Phys. Rev. A* **103**, 032430 (2021).
- Kandala, A. et al. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **549**, 242–246 (2017).

17. Higgott, O., Wang, D. & Brierley, S. Variational quantum computation of excited states. *Quantum* **3**, 156 (2019).
18. Ollitrault, P. J. et al. Quantum equation of motion for computing molecular excitation energies on a noisy quantum processor. *Phys. Rev. Res.* **2**, 043140 (2020).
19. Kandala, A. et al. Error mitigation extends the computational reach of a noisy quantum processor. *Nature* **567**, 491–495 (2019).
20. Bravyi, S., Sheldon, S., Kandala, A., McKay, D. C. & Gambetta, J. M. Mitigating measurement errors in multiqubit experiments. *Phys. Rev. A* **103**, 042605 (2021).
21. Cai, Z., Xu, X. & Benjamin, S. C. Mitigating coherent noise using pauli conjugation. *npj Quantum Inf.* **6**, 1–9 (2020).
22. Grimsley, H. R., Economou, S. E., Barnes, E. & Mayhall, N. J. An adaptive variational algorithm for exact molecular simulations on a quantum computer. *Nat. Commun.* **10**, 1–9 (2019).
23. Tang, H. L. et al. qubit-adapt-vqe: an adaptive algorithm for constructing hardware-efficient ansätze on a quantum processor. *PRX Quantum* **2**, 020310 (2021).
24. McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R. & Neven, H. Barren plateaus in quantum neural network training landscapes. *Nat. Commun.* **9**, 1–6 (2018).
25. Cerezo, M., Sone, A., Volkoff, T., Cincio, L. & Coles, P. J. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nat. Commun.* **12**, 1–12 (2021).
26. Cruise, J. R., Gillespie, N. I. & Reid, B. Practical quantum computing: The value of local computation. Preprint at <https://arxiv.org/abs/2009.08513> (2020).
27. Zhang, D.-B. & Yin, T. Collective optimization for variational quantum eigen-solvers. *Phys. Rev. A* **101**, 032311 (2020).
28. *ibmq_manhattan* (v1.5.1), *ibmq_toronto* (v1.1.4), *ibmq_paris* (v1.6.5), *ibmq_santiago* (v1.3.0), *ibmq_athens* (v1.3.3), *ibmq_valencia* (v1.4.3). IBM Quantum team. <https://quantum-computing.ibm.com> (2020).
29. Peruzzo, A. et al. A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* **5**, 1–7 (2014).
30. Frazier, P. I. A tutorial on Bayesian optimization. Preprint at <https://arxiv.org/abs/1807.02811> (2018).
31. Wigley, P. B. et al. Fast machine-learning online optimization of ultra-cold-atom experiments. *Sci. Rep.* **6**, 1–6 (2016).
32. Nakamura, I., Kanemura, A., Nakaso, T., Yamamoto, R. & Fukuhara, T. Non-standard trajectories found by machine learning for evaporative cooling of 87 rb atoms. *Opt. Express* **27**, 20435–20443 (2019).
33. Mukherjee, R., Xie, H. & Mintert, F. Bayesian optimal control of greenberger-horne-zeilinger states in rydberg lattices. *Phys. Rev. Lett.* **125**, 203603 (2020).
34. Sauvage, F. & Mintert, F. Optimal quantum control with poor statistics. *PRX Quantum* **1**, 020322 (2020).
35. Otterbach, J. et al. Unsupervised machine learning on a hybrid quantum computer. Preprint at <https://arxiv.org/abs/1712.05771> (2017).
36. Zhu, D. et al. Training of quantum circuits on a hybrid quantum computer. *Sci. Adv.* **5**, eaaw9918 (2019).
37. Sung, K. J. et al. Using models to improve optimizers for variational quantum algorithms. *Quantum Sci. Technol.* **5**, 044008 (2020).
38. Ovchinnikov, A., Dmitriev, D., Krivnov, V. Y. & Cheranovskii, V. Antiferromagnetic ising chain in a mixed transverse and longitudinal magnetic field. *Phys. Rev. B* **68**, 214406 (2003).
39. Gray, J. quimb: A python package for quantum information and many-body calculations. *J. Open Source Softw.* **3**, 819 (2018).
40. Abraham, H. et al. Qiskit: an open-source framework for quantum computing. <https://doi.org/10.5281/zenodo.2562110> (2019).
41. Bravyi, S., Gambetta, J. M., Mezzacapo, A. & Temme, K. Tapering off qubits to simulate fermionic hamiltonians. Preprint at <https://arxiv.org/abs/1701.08213> (2017).
42. McClean, J. R. et al. Openfermion: the electronic structure package for quantum computers. *Quantum Sci. Technol.* **5**, 034014 (2020).
43. Bridgeman, J. C. & Chubb, C. T. Hand-waving and interpretive dance: an introductory course on tensor networks. *J. Phys. A* **50**, 223001 (2017).
44. Abadi, M. et al. TensorFlow: large-scale machine learning on heterogeneous systems. <https://www.tensorflow.org/> (2015).
45. Authors. Gpyopt: a Bayesian optimization framework in python. <http://github.com/SheffieldML/GPyOpt> (2016).
46. Snoek, J., Larochelle, H. & Adams, R. P. Practical bayesian optimization of machine learning algorithms. *Adv. Neural Inf. Process Syst.* **25**, 2951–2959 (2012).
47. Shahriari, B., Swersky, K., Wang, Z., Adams, R. P. & De Freitas, N. Taking the human out of the loop: a review of bayesian optimization. *Proc. IEEE* **104**, 148–175 (2015).
48. Bartlett, R. J., Kucharski, S. A. & Noga, J. Alternative coupled-cluster ansätze ii. the unitary coupled-cluster method. *Chem. Phys. Lett.* **155**, 133–140 (1989).
49. Yeter-Aydeniz, K. et al. Benchmarking quantum chemistry computations with variational, imaginary time evolution, and krylov space solver algorithms. *Adv. Quantum Technol.* **4**, 2100012 (2021).
50. Imperial college research computing service. <https://doi.org/10.14469/hpc/2232>.

ACKNOWLEDGEMENTS

We thank Rick Mukherjee for useful discussions on Bayesian optimisation. This work is supported by the Samsung GRC grant, the UK Hub in Quantum Computing and Simulation, part of the UK National Quantum Technologies Programme with funding from UKRI EPSRC grant EP/T001062/1 and the QuantERA ERA-NET Cofund in Quantum Technologies implemented within the European Union's Horizon 2020 Programme. The project TheoryBlind Quantum Control TheBlinQC has received funding from the QuantERA ERA-NET Cofund in Quantum Technologies implemented within the European Unions Horizon 2020 Programme and from EPSRC under the grant EP/R044082/1. F.S. is supported by a studentship in the Quantum Systems Engineering Skills and Training Hub at Imperial College London funded by EPSRC (EP/P510257/1). We acknowledge the use of IBM Quantum services for this work. The views expressed are those of the authors, and do not reflect the official policy or position of IBM or the IBM Quantum team. Numerical simulations in "Testing the effectiveness of information sharing" were carried out on Imperial HPC facilities⁵⁰.

AUTHOR CONTRIBUTIONS

All authors contributed to developing the idea of information sharing. Code implementations of Bayesian optimisation with information sharing were written by K.K., C.S., F.S. and A.S. A.S. developed and implemented the ansatz design scheme. Simulations and experiments were executed by C.S. and K.K. K.K. and C.S. wrote the manuscript. C.S., K.K. and A.S. wrote the supplementary methods. F.M. and F.S. provided expertise on Bayesian optimisation. P.H. and J.K. provided theoretical support on the physical problems used as benchmarks. M.S.K. led and coordinated the project. All authors contributed to revising the manuscript.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Correspondence and requests for materials should be addressed to C.N.S.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021