nature portfolio

Peer Review File

Qubit-excitation-based adaptive variational quantum eigensolver

Reviewers' comments:

Reviewer #1 (Remarks to the Author):

In this work, the authors proposed a method to build the ansatz in VQE iteratively. The new method is based on the adaptive VQE framework and a new building block of the adaptive VQE called qubit excitations was proposed. The new method shows an advantage in reducing CNOT gate count and therefore is more friendly for near-term hardware. This work may be publishable, but there are several places that we hope the authors could fix/explain:

1. This work does not improve the fundamental structure of the previous adaptive VQE method. In our view, it should be called "adaptive VQE with qubit excitation pool" rather than IQEB, which is hard for the readers to identify the content of this work.
2. The authors use numerical simulations on LiH and BeH2 in STO-3G to show the advantage of the new method. However, the difficulty of an electronic structure problem does not depend solely on the number of spin-orbital (qubits) it includes. In the original adaptive VQE work, the numerical simulation was shown on LiH, BeH2, and H6, in which H6 is the most difficult one. In this work, the only simulation on LiH and BeH2 was shown. Though it seems to be large (12, 14 qubits), but it is not enough to prove the general advantage of the new method. At least, a simulation on H6 should be presented.
3. The figures of numerical simulation are hard to read in the manuscript. They are not in vector (SVG) form and the colors are hard to discern. The figures should be improved so that the readers and easily compare the new method and the old ones.
4. Though the authors claim the new method is more CNOT-efficient, in Figure 5 (e) (f) we can see that, if we only care about the CNOT needed for chemical accuracy, qubit-ADAPT-VQE outperforms the new method in all four settings. Does this mean that, in the applications where we only care about chemical accuracy, the new methods are not as CNOT-efficient as qubit-ADAPT-VQE?

Reviewer #2 (Remarks to the Author):

The authors present an iterative variational quantum eigensolver (VQE) algorithm that computes ground state energies and eigenstates of many-body fermionic systems using both quantum and classical processors working in tandem. The algorithm, called IQEB-VQE, is a modified version of the previously introduced ADAPT-VQE algorithm. The main differences are in the choice of operator pool and in an additional, optional algorithmic step in which two operators are added in each iteration of the algorithm. The authors demonstrate that their algorithm can reproduce the exact ground state energies and wave functions of small molecular systems, including BeH2 and LiH, while requiring shorter circuit depths compared to the original ADAPT-VQE (which is often called fermionic-ADAPT-VQE) and the closely related qubit-ADAPT-VQE.

The paper is very well written, and the results are important and timely. To my knowledge, the algorithm presented here outperforms all other existing VQEs in terms of circuit depth, which constitutes an important step toward solving practical problems on near-term quantum devices. I recommend publication provided the following items are addressed:

1. The name of the algorithm is not a great choice. ADAPT-VQE refers to a class of iterative VQEs in

which a wave function ansatz is assembled piece by piece by selecting operators from a pre-defined pool using a criterion based on which of these operators changes the cost function the most. Therefore, IQEB-VQE is a specific example of an ADAPT-VQE algorithm. ADAPT-VQE is by now quite well known in the community, and so using a different name creates a potential for confusion, especially given the wide variety of VQEs that now exist. When a fermionic operator pool is used (the choice originally considered in Ref. [35]), the algorithm is called fermionic-ADAPT-VQE. When a pool based on Pauli strings is used, it is called qubit-ADAPT-VQE. Here, a pool based on linear combinations of Pauli strings is used, suggesting that the algorithm should be called QEB-ADAPT-VQE or something of that sort. Whatever name is used, it should be clear that this is a version of ADAPT, and it should be descriptive of the pool that is used.

2. Following up on point 1 above, it would be clearer if the term “fermionic-ADAPT-VQE” were used in places where the particular operator pool choice considered in Ref. [35] is discussed.

3. At the top of page 5, line 132-133, it is stated that qubit-ADAPT-VQE uses an operator pool based on XY Pauli strings of length 2 or 4. This is inaccurate. Qubit-ADAPT-VQE refers to any pool that consists of individual Pauli strings. Note that a wide variety of such pools were considered in Ref. [36], not just length 2 or 4 Pauli strings involving X and Y.

4. Circuit depth is an important criterion in determining the feasibility of a simulation algorithm designed for NISQ-era devices. However, it is not the only one. Another important criterion is the number of measurements. In Ref. [36], it was shown that there exist “minimal complete pools” consisting of 2N-2 Pauli strings, where N is the number of qubits. This means that qubit-ADAPT-VQE can reduce both the circuit depth and the number of measurements compared to fermionic-ADAPT-VQE. The number of measurements is O(N^5) for qubit-ADAPT-VQE and O(N^8) for fermionic-ADAPT-VQE. The operator pool considered in the present work has N^4 operators, and so the number of measurements is O(N^8) like fermionic-ADAPT-VQE. Are all these operators really necessary? For example, in Ref. [36] it was shown that if one starts with a pool of size N^4, then most of these operators can be discarded without affecting the performance of the algorithm. It would be helpful if the authors could show whether or not this is also true for their pool.

5. In Sec. II.E. and Fig. 3, two “simplified iterative algorithms” are compared. It is not completely clear to me what these algorithms are. Is the one fermionic-ADAPT-VQE and the other IQEB-VQE without the optional 5th step?

6. It is pointed out in Sec. II.E. and Sec. III that the z-strings from the Jordan-Wigner mapping play little or no role in the variational flexibility of the ansatz. It would be worth pointing out that this is consistent with similar findings presented in Ref. [36].

7. A typo: In Eq. (14), the indices on A should be ik.

Reviewer #3 (Remarks to the Author):

The manuscript by Yordanov et al closely follows up on the ADAPT-VQE and qubit-ADAPT-VQE

algorithms. By changing the operator pool and using some of their prior work that introduced efficient circuits for the exponentiated (para)fermionic operators, they show that they can obtain convergence using fewer variational parameters and fewer CNOTs compared to these two prior works on ADAPT.

While this work is very interesting and a valuable contribution to the literature, there are serious issues with the presentation, especially in the abstract and in the name of the algorithm. The presented algorithm is a direct follow-up to ADAPT-VQE (the main difference is the change of pool), yet the authors have chosen a completely different name. This is misleading. The authors should use a name that honestly represents that theirs is a special case of the ADAPT VQE idea. I would recommend something like “symmetrized-qubit-ADAPT”, which much closer represents their contribution and puts it in the context of the work they are building on.

For these reasons, I also found the abstract misleading. The way the work is represented there makes it sound like the authors just discovered ADAPT-VQE. There should be a mention to the prior works, the original ADAPT and qubit-ADAPT VQE in the abstract. The last few sentences of the introduction, lines 64-69, is the information that should appear in the abstract.

On line 59, the authors write that ADAPT-VQE was “one of the first” iterative VQE algorithms, while it was, in fact, the first.

It is not clear what circuits are used for the fermionic operators in the case of the fermionic ADAPT-VQE. If not done already, the authors should use their prior construction (Ref. 46) for a fair comparison. Perhaps this is already mentioned, but I couldn’t find the information.

It is not clear how many operators are typically chosen out of the pool because of their high gradients and taken to the full VQE step to down select. This information should be provided, both as a ballpark, but also in the specific cases studied.

A related point is that in the comparison of the three algorithms (Fig. 5 and the surrounding discussion), the metrics used are: variational parameters, iterations, and number of CNOTs. However, this fails to take into account the fact that the authors’ new algorithm includes multiple VQE subroutines (instead of just one) at every iteration. This increases the number of measurements and optimizations, making the overall runtime potentially longer (this depends on how much their approach helps with convergence, see next point). It would be fair to also include a metric related to this in the discussion and possibly in the plots.

Another related point is that the multiple VQE subroutines might not be necessary. It would be interesting to see how the algorithm performs if one of the operators with the highest gradient is chosen at random so that only one VQE subroutine is needed (as in the prior versions of ADAPT). Does that affect the convergence significantly?

Line 325: NISQ is misspelled.

To conclude, this work is a valuable contribution to the literature and should be of interest to the community working on VQEs. However, the authors should strive to be intellectually honest when presenting their results and contributions, see comments above. After the suggested changes are

implemented and my questions addressed, I can recommend publication.

Rebuttal Letter

Responses to reviewers' comments and suggestions:

Reviewer 1

1. This work does not improve the fundamental structure of the previous adaptive VQE method. In our view, it should be called "adaptive VQE with qubit excitation pool" rather than IQEB, which is hard for the readers to identify the content of this work.

We acknowledge that the previous name of our protocol did not properly emphasize its relation to the ADAPT-VQE. Therefore, we renamed our protocol to "Qubit-excitation-based adaptive variational quantum eigensolver", or "QEB-ADAPT-VQE" for short, as suggested by Reviewer 2.

2. The authors use numerical simulations on LiH and BeH2 in STO-3G to show the advantage of the new method. However, the difficulty of an electronic structure problem does not depend solely on the number of spin-orbital (qubits) it includes. In the original adaptive VQE work, the numerical simulation was shown on LiH, BeH2, and H6, in which H6 is the most difficult one. In this work, the only simulation on LiH and BeH2 was shown. Though it seems to be large (12, 14 qubits), but it is not enough to prove the general advantage of the new method. At least, a simulation on H6 should be presented.

We thank Reviewer 1 about this suggestion. We have added simulation results for H6 in all Sections of the paper.

3. The figures of numerical simulation are hard to read in the manuscript. They are not in vector (SVG) form and the colors are hard to discern. The figures should be improved so that the readers and easily compare the new method and the old ones.

We thank Reviewer 1 about pointing out this problem. We have now uploaded figures with higher resolution.

4. Though the authors claim the new method is more CNOT-efficient, in Figure 5 (e) (f) we can see that, if we only care about the CNOT needed for chemical accuracy, qubit-ADAPT-VQE outperforms the new method in all four settings. Does this mean that, in the applications where we only care about chemical accuracy, the new methods are not as CNOT-efficient as qubit-ADAPT-VQE?

We agree with Reviewer 1, that in practice one would be interested to achieve chemical accuracy only. In the presented simulation, the qubit-ADAPT-VQE is indeed more CNOT-efficient than the QEB-ADAPT-VQE in achieving chemical accuracy for LiH and BeH2. Although our results are not sufficient to draw a general conclusion, they indicated that the circuit-efficiency of the QEB-ADAPT-VQE becomes more significant for larger ansatz circuits. In the case of H6, which requires the largest ansatz, the QEB-ADAPT-VQE is more CNOT-efficient than the qubit-ADAPT-VQE at chemical accuracy. Therefore, we anticipate that for larger molecular systems, which would require larger ansatz circuits, the QEB-ADAPT-VQE to be more circuit-efficient than the qubit-ADAPT-VQE even at chemical accuracy level. We now comment on this in the Conclusion of the manuscript.

Reviewer 2

1. The name of the algorithm is not a great choice. ADAPT-VQE refers to a class of iterative VQEs in which a wave function ansatz is assembled piece by piece by selecting operators from a pre-defined pool using a criterion based on which of these operators changes the cost function the most. Therefore, IQEB-VQE is a specific example of an ADAPT-VQE algorithm. ADAPT-VQE is by now quite well known in the community, and so using a different name creates a potential for confusion, especially given the wide variety of VQEs that now exist. When a fermionic operator pool is used (the choice originally considered in Ref. [35]), the algorithm is called fermionic-ADAPT-VQE. When a pool based on Pauli strings is used, it is called qubit-ADAPT-VQE. Here, a pool based on linear combinations of Pauli strings is used, suggesting that the algorithm should be called QEB-ADAPT-VQE or something of that sort. Whatever name is used, it should be clear that this is a version of ADAPT, and it should be descriptive of the pool that is used.

We acknowledge that the previous name of our protocol did not properly emphasize its relation to the ADAPT-VQE. Therefore, we renamed our protocol to "Qubit-excitation-based adaptive variational quantum eigensolver", or "QEB-ADAPT-VQE" for short, as suggested by Reviewer 2.

2. Following up on point 1 above, it would be clearer if the term "fermionic-ADAPT-VQE" were used in places where the particular operator pool choice considered in Ref. [35] is discussed.

We thank Reviewer 2 for this suggestion. We now refer to the original ADAPT-VQE as fermionic-ADAPT-VQE.

3. At the top of page 5, line 132-133, it is stated that qubit-ADAPT-VQE uses an operator pool based on XY Pauli strings of length 2 or 4. This is inaccurate. Qubit-ADAPT-VQE refers to any pool that consists of individual Pauli strings. Note that a wide variety of such pools were considered in Ref. [36], not just length 2 or 4 Pauli strings involving X and Y.

We thank Reviewer 2 for pointing out this inaccuracy. In Section II H of the manuscript, we now elaborate that we implement the qubit-ADAPT-VQE with a pool of odd XY-Pauli strings of length 2 and 4, for the purpose of obtaining the plots in Fig. 6. The evolutions of Pauli strings in this pool can be combined to obtain the qubit excitation evolutions in the pool of the QEB-ADAPT-VQE. Hence, using this pool for the qubit-ADAPT-VQE, we consider the comparison between the two protocols to be fair.

4. Circuit depth is an important criterion in determining the feasibility of a simulation algorithm designed for NISQ-era devices. However, it is not the only one. Another important criterion is the number of measurements. In Ref. [36], it was shown that there exist "minimal complete pools" consisting of 2N-2 Pauli strings, where N is the number of qubits. This means that qubit-ADAPT-VQE can reduce both the circuit depth and the number of measurements compared to fermionic-ADAPT-VQE. The number of measurements is O(N^5) for qubit-ADAPT-VQE and O(N^8) for fermionic-ADAPT-VQE. The operator pool considered in the present work has N^4 operators, and so the number of measurements is O(N^8) like fermionic-ADAPT-VQE. Are all these operators really necessary? For example, in Ref. [36] it was shown that if one starts with a pool of size N^4, then most of these operators can be discarded without affecting the performance of the algorithm. It would be helpful if the authors could show whether or not this is also true for their pool.

This is indeed an interesting question. We believe that a minimal complete pool of O(N) qubit excitations does exist. Acting upon the HF state, such a pool should be able to span the whole real part of the Hilbert subspace of Hamming weight "*h*", where "*h*" is equal to the number of electrons. However, for the moment we have not been able to derive an analytical proof that such a reduced pool of qubit excitation evolutions exists. This is something we currently work on and hope to present in future.

Regarding, the QEB-ADAPT-VQE. we are not sure how useful a reduced pool of qubit excitation evolutions would be. First, a smaller pool will result in a slower and less efficient ansatz construction.
Second, as we discuss in Section IV of the Supplementary information, the total complexity of the QEB-ADAPT-VQE might actually be dominated by the complexity of running the VQE at each iteration, which can be as large as O(N^12). Therefore, using a reduced ansatz element pool will decrease the number of measurements required to evaluate the energy gradients at each iteration from O(N^8) to O(N^5), but will not change the overall complexity of the protocol. We now comment on this in the Conclusion of the paper.

5. In Sec. II.E. and Fig. 3, two "simplified iterative algorithms" are compared. It is not completely clear to me what these algorithms are. Is the one fermionic-ADAPT-VQE and the other IQEB-VQE without the optional 5th step?

We thank Reviewer 2 for his/her comment. We now elaborate, in Section II F, that we compare the QEB-ADAPT-VQE for n=1 and step 5 not performed, and the fermionic-ADAPT-VQE using an ansatz element pool of individual (non-spin-complement) fermionic excitation evolutions.

6. It is pointed out in Sec. II.E. and Sec. III that the z-strings from the Jordan-Wigner mapping play little or no role in the variational flexibility of the ansatz. It would be worth pointing out that this is consistent with similar findings presented in Ref. [36].

As suggested by Reviewer 2, in the conclusion of our paper, we point out that similar findings were presented in Ref. [36].

7. A typo: In Eq. (14), the indices on A should be ik.

We thank the Reviewer about pointing out the typo. It is now fixed.

Reviewer 3

1. While this work is very interesting and a valuable contribution to the literature, there are serious issues with the presentation, especially in the abstract and in the name of the algorithm. The presented algorithm is a direct follow-up to ADAPT-VQE (the main difference is the change of pool), yet the authors have chosen a completely different name. This is misleading. The authors should use a name that honestly represents that theirs is a special case of the ADAPT VQE idea. I would recommend something like "symmetrized-qubit-ADAPT", which much closer represents their contribution and puts it in the context of the work they are building on.

We acknowledge that the previous name of our protocol did not properly emphasize its relation to the ADAPT-VQE. Therefore, we renamed our protocol to "Qubit-excitation-based adaptive variational quantum eigensolver", or "QEB-ADAPT-VQE" for short, as suggested by Reviewer 2.

2. For these reasons, I also found the abstract misleading. The way the work is represented there makes it sound like the authors just discovered ADAPT-VQE. There should be a mention to the prior works, the original ADAPT and qubit-ADAPT VQE in the abstract. The last few sentences of the introduction, lines 64-69, is the information that should appear in the abstract.

Based on the suggestions of the Reviewers, we rewrote parts of the Abstract and the Introduction to better emphasize the relation of the QEB-ADAPT-VQE to the ADAPT-VQE.

3. On line 59, the authors write that ADAPT-VQE was “one of the first” iterative VQE algorithms, while it was, in fact, the first.

We thank Reviewer 3 for pointing out this inaccuracy. This part of the Introduction has been rewritten.

4. It is not clear what circuits are used for the fermionic operators in the case of the fermionic ADAPT-VQE. If not done already, the authors should use their prior construction (Ref. 46) for a fair comparison. Perhaps this is already mentioned, but I couldn’t find the information.

We thank Reviewer 3 about pointing out this. We mention in Section II C and the caption of Fig. 6 that for fermionic excitation evolutions, we consider the use of the CNOT-optimized circuits, which we previously derived in Ref. [46]. Therefore, for both the QEB-ADAPT-VQE and the fermionic-ADAPT-VQE we consider the same circuits optimizations.

5. It is not clear how many operators are typically chosen out of the pool because of their high gradients and taken to the full VQE step to down select. This information should be provided, both as a ballpark, but also in the specific cases studied.

We thank Reviewer 3 about his/her comment. For the majority of our results, we chose only one operator (n=1). The only exceptions are the results presented in Fig. 5. In those cases, the QEB-ADAPT-VQE is performed with n=10. Previously, this information was only given, cursorily, in the main body of the text. Now, we include this information in the captions of the corresponding figures as well.

6. A related point is that in the comparison of the three algorithms (Fig. 5 and the surrounding discussion), the metrics used are: variational parameters, iterations, and number of CNOTs. However, this fails to take into account the fact that the authors’ new algorithm includes multiple VQE subroutines (instead of just one) at every iteration. This increases the number of measurements and optimizations, making the overall runtime potentially longer (this depends on how much their approach helps with convergence, see next point). It would be fair to also include a metric related to this in the discussion and possibly in the plots.

We thank Reviewer 3 for his/her comment. We realized that in the previous version of the paper we actually compared the QEB-ADAPT-VQE for n=10, to the other two protocols, which was not a fair comparison. Now, when comparing the QEB-ADAPT-VQE to the fermionic-ADAPT-VQE and the qubit-ADAPT-VQE, we perform the QEB-ADAPT-VQE for n=1 (so just a single VQE subroutine is run at every iteration). Because of the valuable comment made by the referee, we believe that our present manuscript gives a fair comparison.

7. Another related point is that the multiple VQE subroutines might not be necessary. It would be interesting to see how the algorithm performs if one of the operators with the highest gradient is chosen at random so that only one VQE subroutine is needed (as in the prior versions of ADAPT). Does that affect the convergence significantly?

As suggested by Reviewer 3, we added a section in the Supplementary information, where we study the performance of the QEB-ADAPT-VQE for different values of n. For n=10, a CNOT count reduction of about 15%-25% can be achieved on average for the three studied molecules.

8. Line 325: NISQ is misspelled.

We thank the Reviewer for pointing out the typo. It is now fixed.

REVIEWERS' COMMENTS:

Reviewer #1 (Remarks to the Author):

The authors have addressed our comments. The manuscript is now ready for publication.

Reviewer #2 (Remarks to the Author):

The authors have responded thoroughly to all of my comments and questions and updated the manuscript accordingly with substantial revisions. In my opinion, they have also satisfactorily addressed all the issues raised by the other referees as well. I recommend that the manuscript now be published in its current form.

Reviewer #3 (Remarks to the Author):

The authors have addressed in depth the comments and concerns of all three reviewers. I find that the manuscript has improved significantly. I now recommend publication in Communications Physics.