Peer Review Information

Journal: Nature Computational Science
Manuscript Title: Provable Bounds for Noise-Free Expectation Values Computed from Noisy Samples
Corresponding author name(s): Woerner

Editorial Notes:

None

Reviewer Comments & Decisions:

Decision Letter, initial version:

Dear Dr Woerner,

Your manuscript "Provable Bounds for Noise-Free Expectation Values Computed from Noisy Samples" has now been seen by 2 referees, whose comments are appended below.

Please accept our sincerest apologies for the length of time your manuscript has been under consideration at our journal - one referee was eventually unable to review the manuscript, despite our persistent efforts to reduce the delays. We thank you very much for your patience during what has been an uncharacteristically long peer-review round.

You will see that while they find your work of interest, they have raised points that need to be addressed before we can make a decision on publication.

The referees' reports seem to be quite clear. Naturally, we will need you to address *all* of the points raised.

While we ask you to address all of the points raised, the following points need to be substantially worked on:
- Provide the experiments suggested by our referees to broaden the scope of your study, including different noise models, algorithms other than QAOA

- Provide more experiments to better support the claims and to show the advantage, as suggested from our referees
- Include more discussions about technical details, limitations and future research directions

Please use the following link to submit your revised manuscript and a point-by-point response to the referees' comments (which should be in a separate document to any cover letter):

[REDACTED]

** This url links to your confidential homepage and associated information about manuscripts you may have submitted or be reviewing for us. If you wish to forward this e-mail to co-authors, please delete this link to your homepage first. **

To aid in the review process, we would appreciate it if you could also provide a copy of your manuscript files that indicates your revisions by making use of Track Changes or similar mark-up tools. Please also ensure that all correspondence is marked with your Nature Computational Science reference number in the subject line.

In addition, please make sure to upload a Word Document or LaTeX version of your text, to assist us in the editorial stage.

If you have any issues when updating your Code Ocean capsule during the revision process, please email the Code Ocean support team Cc'ing me.

To improve transparency in authorship, we request that all authors identified as 'corresponding author' on published papers create and link their Open Researcher and Contributor Identifier (ORCID) with their account on the Manuscript Tracking System (MTS), prior to acceptance. ORCID helps the scientific community achieve unambiguous attribution of all scholarly contributions. You can create and link your ORCID from the home page of the MTS by clicking on 'Modify my Springer Nature account'. For more information please visit please visit www.springernature.com/orcid.

We hope to receive your revised paper within three weeks. If you cannot send it within this time, please let us know.

We look forward to hearing from you soon.

Best regards,

Jie Pan, Ph.D.

Senior Editor
Nature Computational Science

Reviewers comments:

Reviewer #1 (Remarks to the Author):

This paper studies noises in quantum computers, and introduces a quantification method to provide bounds on expectation values. The applications of this method include optimization-related and fidelity-related tasks, and QAOA experiments have been demonstrated on IBM superconducting quantum devices.

This manuscript is in general well-written and organized. After introducing the theoretical framework, the following applications exhibit practical utility. We evaluate positively on this work, but some major and minor points should be properly addressed before we make final recommendations.

Major Comments

1. The noise model considered in this work is a specific and simple one. This model allows for analytical derivations, but how can this be generalized to other noise models?

2. Why are the experiments put in a QAOA context? In the two experiments, the first one is classically optimized, and the second one is done by the parameter transfer trick. They are actually not QAOA experiments and QAOA optimization is not involved. Why not simply do an expectation value evaluation task that is less misleading and captures the spirit?

Other Comments

1. The introduction of the Quantum Alternating Operator Ansatz at the end of IV.A is rather brief. For a note, it is enough. For a paper, it should provide more details for the general audience.

2. In the conclusion, the authors mention "The methodologies developed in this paper 'can be adapted to account for SPAM errors', either by increasing sampling overhead or applying other mitigation techniques, like statistical readout error mitigation". The authors make this judgment about SPAM errors but then leave it "for future research", which is not appropriate. Such words should be supported by proofs/data.

Reviewer #2 (Remarks to the Author):

The paper studies the effect of noise on sampling from noisy quantum devices. This is a well-studied topic, and many existing methods for mitigating noise exist. Examples include zero noise extrapolation and probabilistic error cancellation. These error mitigation techniques allow one to estimate expectation values accurately using noisy quantum devices with the cost of having a larger measurement overhead. However, mitigating noise in sampling remains a challenge.

The contributions of this work are the following:

(a) The authors presented a simple observation that if a bitstring x has a probability px to be sampled in the noiseless case, then x has a probability of at least px/sqrt {gamma} to be sampled in the noisy case with gamma being a measure of the noise strength (larger gamma is higher noise). This observation relies on a Pauli-Lindblad noise model recently introduced, which accurately describes various noisy quantum devices.

(b) Using the simple observation above, the authors proved a relatively straightforward lemma that a very simple previously proposed heuristic estimate, "CONDITIONAL VALUE-AT-RISK," can be used to obtain provable bounds on the true expectation values. However, how good the bounds are in capturing the expectation values can only be justified heuristically in experiments (if the bounds are very loose, then they are not useful).

(c) The authors experimentally demonstrated the use of these two observations (the one about sampling and the one about bounding expectation values) in approximate quantum optimization problems.

Given the assumption made in this work, the two main analytical contributions are very easy to see and show. There are not too many surprises from the mathematical side in deducing the two observations from the noise model assumption.

While the main results in the paper are simple to prove, I think an important part of this work is in "discovering" the implications of the assumption about the noise model. The noise model is a reasonable model proposed in an earlier work. The work shows that this reasonable assumption actually has relatively nice consequences. Overall, I find the observations presented in this work to be interesting from a practical point of view.

I have a few suggestions that could potentially make this work better:

(1) The authors should provide a better justification for using the bounds provided by CVaR compared to PEC and ZNE. The introduction is missing this important rationale for why we want to use CVaR, which only provides bounds, instead of PEC, which provides accurate estimates.

(2) The experiments should also be expanded to support the advantage of considering the noise mitigation theory via CVaR compared to PEC and ZNE.

(3) The experiments seem to focus solely on QAOA, but the text suggests other applications (such as QSVMs and VarQTE). I think having a few numerics on these other applications will be very helpful.

(4) The conclusion section did not provide sufficient future prospects. Currently, the conclusion only mentions that SPAM errors can also be studied, which seems like a rather limited prospect. I would appreciate more discussions about how CVaR and the observations may be further enhanced as well as more diverse future directions.

Reviewer #2 (Remarks on code availability):

The code provides a minimal yet useful README file for running the code. I was able to install and run the code provided in the Jupyter notebook. I did not go through all outputs of the code in great detail. But I have verified that many of the key findings stated in the numerical experiments can be generated from the code.

Author Rebuttal to Initial comments

Reply to the Editor's and Reviewers' Comments (NATCOMPUTSCI-24-0665A)

Dear Jie Pan,

Thanks a lot for for the constructive feedback on our work and for the invitation to submit a revised version of our manuscript. We carefully addressed all the reviewers' comments as well as the points you highlighted.

In particular, we implemented the following changes to our work:

- We added additional experiments demonstrating on 50 as well as 100 qubits how our theory applies to fidelity estimation, a core building block for Quantum Support Vector Machines and Variational Quantum Time Evolution. These experiments highlight the close agreement of our theory with the experimental results as well as the practical relevance of our results. We conclude these experiments by providing more intuition about when we expect the bounds to be tight and when not.
- We (re-)ran all included experiments using measurement error mitigation. Further, we discuss in more detail alternative approaches to mitigate SPAM errors in future work.
- We adjusted the main text to highlight that our theory also applies to more general noise models, and that the mentioned Pauli-Lindblad noise model is only used for ease of presentation and to relate to existing literature on error mitigation.
- We expanded the introduction as well as conclusion to better put our work into context and provide more suggestions for future research.

Below, you will find the reviewers' comments as well as our detailed replies (in blue). We believe the revised manuscript has significantly improved and hope you find it suitable for publication in *Nature Computational Science*. We are looking forward to hearing back from you.

Best regards

Stefan Woerner, on behalf of the authors

Reviewer #1 (Remarks to the Author):

Dear Reviewer, thanks a lot for your constructive feedback on our manuscript. We addressed all your comments, which helped us to significantly improve our paper. Please find the detailed answers to your comments below.

This paper studies noises in quantum computers, and introduces a quantification method to provide bounds on expectation values. The applications of this method include optimization-related and fidelity-related tasks, and QAOA experiments have been demonstrated on IBM superconducting quantum devices.

This manuscript is in general well-written and organized. After introducing the theoretical framework, the following applications exhibit practical utility. We evaluate positively on this work, but some major and minor points should be properly addressed before we make final recommendations.

Thank you for the positive evaluation of our work and the constructive feedback.

Major Comments

1. The noise model considered in this work is a specific and simple one. This model allows for analytical derivations, but how can this be generalized to other noise models?

Our results hold for more general noise models. We focus on the Pauli-Lindblad noise model for ease of presentation, because it has been empirically shown to accurately describe many quantum devices, and to be able to link our results to the related literature on Probabilistic Error Cancelation (PEC). We revised the manuscript to highlight the more general applicability of our results.

2. Why are the experiments put in a QAOA context? In the two experiments, the first one is classically optimized, and the second one is done by the parameter transfer trick. They are actually not QAOA experiments and QAOA optimization is not involved. Why not simply do an expectation value evaluation task that is less misleading and captures the spirit?

We re-organized the "Experiments" section and start with a new experiment on fidelity estimation to demonstrate our bounds in that context. Fidelity estimation is a key building block for Quantum Support Vector Machines and Variational Quantum Time Evolution. The experiments show a nice agreement between our theory and the experimental results.

Further, we expanded the discussion on QAOA. There is an increasing number of papers proposing alternative ways to select QAOA parameters, either by training them classically, transferring parameters, or applying other rules, e.g., inspired by quantum annealing or infinite problem-size limits. In this spirit, our results show how to compensate noise for sampling from a quantum computer and with the CVaR we validate that the theory holds.

Other Comments

1. The introduction of the Quantum Alternating Operator Ansatz at the end of IV.A is rather brief. For a note, it is enough. For a paper, it should provide more details for the general audience.

We expanded our introduction of the Quantum Alternating Operator Ansatz and how filtering / post-selection could be used to strengthen our CVaR bounds on the corresponding expectation values.

2. In the conclusion, the authors mention "The methodologies developed in this paper 'can be adapted to account for SPAM errors', either by increasing sampling overhead or applying other mitigation techniques, like statistical readout error mitigation". The authors make this judgment about SPAM errors but then leave it "for future research", which is not appropriate. Such words should be supported by proofs/data.

We re-ran all our experiments using M3 measurement error mitigation. Further, we provide more details and references on defining of SPAM errors as Pauli noise, which allows to apply our theory. This leads to a more consistent treatment of SPAM errors throughout the whole manuscript.

Reviewer #2 (Remarks to the Author):

Dear Reviewer, thanks a lot for your constructive feedback on our manuscript. It helped us a lot to significantly improve the revised version. Please find the detailed answers to your comments below.

The paper studies the effect of noise on sampling from noisy quantum devices. This is a well-studied topic, and many existing methods for mitigating noise exist. Examples include zero noise extrapolation and probabilistic error cancellation. These error mitigation techniques allow one to estimate expectation values accurately using noisy quantum devices with the cost of having a larger measurement overhead. However, mitigating noise in sampling remains a challenge.

The contributions of this work are the following:

(a) The authors presented a simple observation that if a bitstring x has a probability px to be sampled in the noiseless case, then x has a probability of at least px/sqrt{gamma} to be sampled in the noisy case with gamma being a measure of the noise strength (larger gamma is higher noise). This observation relies on a Pauli-Lindblad noise model recently introduced, which accurately describes various noisy quantum devices.

(b) Using the simple observation above, the authors proved a relatively straightforward lemma that a very simple previously proposed heuristic estimate, "CONDITIONAL VALUE-AT-RISK," can be used to obtain provable bounds on the true expectation values. However, how good the bounds are in capturing the expectation values can only be justified heuristically in experiments (if the bounds are very loose, then they are not useful).

(c) The authors experimentally demonstrated the use of these two observations (the one about sampling and the one about bounding expectation values) in approximate quantum optimization problems.

Given the assumption made in this work, the two main analytical contributions are very easy to see and show. There are not too many surprises from the mathematical side in deducing the two observations from the noise model assumption.

While the main results in the paper are simple to prove, I think an important part of this work is in "discovering" the implications of the assumption about the noise model. The noise model is a reasonable model proposed in an earlier work. The work shows that this reasonable assumption actually has relatively nice consequences. Overall, I find the observations presented in this work to be interesting from a practical point of view.

Thank you for the positive evaluation of our work and the constructive feedback.

Please note that most of our results hold also under more general noise models. We focused on the Pauli-Lindblad noise model for ease of presentation and to relate to existing literature. We adjusted the manuscript accordingly and highlight the more general applicability of our results.

While Lemma 1 can be proven easily, we want to highlight also Lemma 3 (in Appendix C), which is a bit more advanced to prove, implying a quartic reduction in variance amplification for the CVaR bounds compared to PEC for expectation values.

I have a few suggestions that could potentially make this work better:

(1) The authors should provide a better justification for using the bounds provided by CVaR compared to PEC and ZNE. The introduction is missing this important rationale for why we want to use CVaR, which only provides bounds, instead of PEC, which provides accurate estimates.

We added a justification to the introduction, summarizing that CVaR does not require expensive learning of a noise model, needs less restrictive assumptions on the noise model, and implies a significantly lower sampling overhead compared to PEC. Compared to ZNE, similar advantages hold, and in particular, ZNE is usually only a heuristic without the theoretical guarantees of PEC and CVaR.

(2) The experiments should also be expanded to support the advantage of considering the noise mitigation theory via CVaR compared to PEC and ZNE.

We added a detailed discussion in the introduction to highlight the advantages of CVaR over PEC and ZNE. Further, we added a new experiment on fidelity estimation (also see next comment) that nicely highlights how powerful the CVaR bounds can be with relatively few shots (compared to, e.g., PEC).

(3) The experiments seem to focus solely on QAOA, but the text suggests other applications (such as QSVMs and VarQTE). I think having a few numerics on these other applications will be very helpful.

We re-organized the "Experiments" section as follows: We start with a new experiment on fidelity estimation, which nicely demonstrates the agreement of our theory with the experimental results. This is the key building block for the other mentioned algorithms. Then we discuss the QAOA experiments, that we all re-ran using measurement error mitigation.

(4) The conclusion section did not provide sufficient future prospects. Currently, the conclusion only mentions that SPAM errors can also be studied, which seems like a rather limited prospect. I would appreciate more discussions about how CVaR and the observations may be further enhanced as well as more diverse future directions.

We add more details on handling SPAM errors to main text and conclusion. Further, we propose using the demonstrated fidelity estimation in the discussed algorithms, as well as training QAOA circuits on hardware using CVaR. Further, we discuss the opportunities that filtering / post-selection offers to further strengthen the CVaR bounds, either by leveraging natural symmetries or by introducing certain properties through careful modeling of problems.

Reviewer #2 (Remarks on code availability):

The code provides a minimal yet useful README file for running the code. I was able to install and run the code provided in the Jupyter notebook. I did not go through all outputs of the code in great detail. But I have verified that many of the key findings stated in the numerical experiments can be generated from the code.

Thank you again for the positive and constructive feedback, which helped to significantly improve our manuscript.

Decision Letter, first revision:

26th August 2024

Dear Dr. Woerner,

Thank you for submitting your revised manuscript "Provable Bounds for Noise-Free Expectation Values Computed from Noisy Samples" (NATCOMPUTSCI-24-0665B). It has now been seen by the original referees and their comments are below. The reviewers find that the paper has improved in revision, and therefore we'll be happy in principle to publish it in Nature Computational Science, pending minor revisions to satisfy the referees' final requests and to comply with our editorial and formatting guidelines.

We are now performing detailed checks on your paper and will send you a checklist detailing our editorial and formatting requirements in about two weeks. Please do not upload the final materials and make any revisions until you receive this additional information from us.

TRANSPARENT PEER REVIEW
Nature Computational Science offers a transparent peer review option for original research manuscripts. We encourage increased transparency in peer review by publishing the reviewer comments, author rebuttal letters and editorial decision letters if the authors agree. Such peer review material is made available as a supplementary peer review file. Please remember to choose, using the manuscript system, whether or not you want to participate in transparent peer review.
Please note: we allow redactions to authors' rebuttal and reviewer comments in the interest of confidentiality. If you are concerned about the release of confidential data, please let us know specifically what information you would like to have removed. Please note that we cannot incorporate redactions for any other reasons. Reviewer names will be published in the peer review files if the reviewer signed the comments to authors, or if reviewers explicitly agree to release their name. For more information, please refer to our FAQ page.

Thank you again for your interest in Nature Computational Science. Please do not hesitate to contact me if you have any questions.

Sincerely,

Jie Pan, Ph.D.
Senior Editor
Nature Computational Science

ORCID
IMPORTANT: Non-corresponding authors do not have to link their ORCIDs but are encouraged to do so. Please note that it will not be possible to add/modify ORCIDs at proof. Thus, please let your co-authors know that if they wish to have their ORCID added to the paper they must follow the procedure described in the following link prior to acceptance:
https://www.springernature.com/gp/researchers/orcid/orcid-for-nature-research

Reviewer #1 (Remarks to the Author):

I went through both referees' reports, the revised manuscript, and the authors' reply. It is clear that the authors have made great efforts to address both referees' comments and suggestions. They have added substantial new results/discussions and the revised manuscript is significantly improved. In particular, they added additional experiments to support their claims and clarified some confusing points. As far as I am concerned, the authors have satisfactorily addressed all my concerns/suggests raised in the first round. With this substantially improved manuscript, I am convinced that this work might be suitable for publication in Nature Computational Science and I am happy to recommend its acceptance.

Reviewer #2 (Remarks to the Author):

I would like to thank the authors for taking a substantive revision to the manuscript.

The authors have addressed all points raised in my previous review. The newly added experiments on fidelity estimation very nicely demonstrate the advantage of CVaR over existing methods, such as PEC. The additional texts provide better explanations for the usefulness of CVaR.

The revised manuscript is much better than the original submission.
I would hence recommend acceptance to Nature Computational Science.

Final Decision Letter:

Dear Dr Woerner,

We are pleased to inform you that your Article "Provable Bounds for Noise-Free Expectation Values Computed from Noisy Samples" has now been accepted for publication in Nature Computational Science.

Once your manuscript is typeset, you will receive an email with a link to choose the appropriate publishing options for your paper and our Author Services team will be in touch regarding any additional information that may be required.

Please note that Nature Computational Science is a Transformative Journal (TJ). Authors may publish their research with us through the traditional subscription access route or make their paper immediately open access through payment of an article-processing charge (APC). Authors will not be required to make a final decision about access to their article until it has been accepted. Find out more about Transformative Journals

Authors may need to take specific actions to achieve compliance with funder and institutional open access mandates. If your research is supported by a funder that requires immediate open access (e.g. according to Plan S principles) then you should select the gold OA route, and we will direct you to the compliant route where possible. For authors selecting the subscription publication route, the journal's standard licensing terms will need to be accepted, including self-archiving policies. Those licensing terms will supersede any other terms that the author or any third party may assert apply to any version of the manuscript.

If you have any questions about our publishing options, costs, Open Access requirements, or our legal forms, please contact ASJournals@springernature.com

Acceptance of your manuscript is conditional on all authors' agreement with our publication policies (see https://www.nature.com/natcomputsci/for-authors). In particular your manuscript must not be published elsewhere and there must be no announcement of the work to any media outlet until the publication date (the day on which it is uploaded onto our web site).

Before your manuscript is typeset, we will edit the text to ensure it is intelligible to our wide readership and conforms to house style. We look particularly carefully at the titles of all papers to ensure that they are relatively brief and understandable.

Once your manuscript is typeset, you will receive a link to your electronic proof via email with a request to make any corrections within 48 hours. If, when you receive your proof, you cannot meet this deadline, please inform us at rjsproduction@springernature.com immediately.

If you have queries at any point during the production process then please contact the production team at rjsproduction@springernature.com.

You may wish to make your media relations office aware of your accepted publication, in case they consider it appropriate to organize some internal or external publicity. Once your paper has been scheduled you will receive an email confirming the publication details. This is normally 3-4 working days in advance of publication. If you need additional notice of the date and time of publication, please let the production team know when you receive the proof of your article to ensure there is sufficient time to coordinate. Further information on our embargo policies can be found here: https://www.nature.com/authors/policies/embargo.html

An online order form for reprints of your paper is available at https://www.nature.com/reprints/author-reprints.html. All co-authors, authors' institutions and authors' funding agencies can order reprints using the form appropriate to their geographical region.

We welcome the submission of potential cover material (including a short caption of around 40 words) related to your manuscript; suggestions should be sent to Nature Computational Science as electronic files (the image should be 300 dpi at 210 x 297 mm in either TIFF or JPEG format). We also welcome suggestions for the Hero Image, which appears at the top of our home page; these should be 72 dpi at 1400 x 400 pixels in JPEG format. Please note that such pictures should be selected more for their aesthetic appeal than for their scientific content, and that colour images work better than black and white or grayscale images. Please do not try to design a cover with the Nature Computational Science logo etc., and please do not submit composites of images related to your work. I am sure you will understand that we cannot make any promise as to whether any of your suggestions might be selected for the cover of the journal.

You can now use a single sign-on for all your accounts, view the status of all your manuscript submissions and reviews, access usage statistics for your published articles and download a record of your refereeing activity for the Nature journals.

To assist our authors in disseminating their research to the broader community, our SharedIt initiative provides you with a unique shareable link that will allow anyone (with or without a subscription) to read the published article. Recipients of the link with a subscription will also be able to download and print the PDF.

As soon as your article is published, you will receive an automated email with your shareable link.

We look forward to publishing your paper.

Best regards,

Jie Pan, Ph.D.
Senior Editor
Nature Computational Science

P.S. Click on the following link if you would like to recommend Nature Computational Science to your librarian: https://www.springernature.com/gp/librarians/recommend-to-your-library

** Visit the Springer Nature Editorial and Publishing website at www.springernature.com/editorial-and-publishing-jobs for more information about our career opportunities. If you have any questions please click here.**