

# A coherent approach to quantum-classical optimization

Corresponding Author: Mr Jordi Riu Vicente

**This file contains all reviewer reports in order by version, followed by all author rebuttals in order by version.**

Version 0:

Reviewer comments:

Reviewer #1

(Remarks to the Author)

The authors propose a classical-quantum framework variant for binary optimization problems. Extending previously suggested frameworks involving classical tensor networks and variational quantum algorithms, the key insight is that the quantum component (the QAOA algorithm) applied to a binary optimization problem greatly benefits from an initialization that is both low cost and high entropy in computational basis. This requires that the classical tensor network, here a matrix product state (MPS), uses a different training protocol that optimizes for both of those properties.

I find this work very insightful, and I have almost no comments to the first half of the manuscript. The finding that high entropy initializations are vital for the QAOA algorithm is very relevant due to the nature of the classical optimization problems where MPS are likely to greedily optimize for the lowest energy basis state. The one issue I have with this insight, also in Appendix A, is that it neglects the possibility to early-stop the MPS training when entropy or the bond dimension starts to drop. A comparison to such an approach would be very nice.

The contextualization and motivation of the imaginary time evolution approach via the Helmholtz free energy cost function is also very nice. Though when it comes to actually applying the imaginary time evolution, the very rough approximation via the first order method is not appropriately studied. For example, it is not clear how close to the Boltzmann limit of entropy the evolution happens, or how well the solution could have been classically given this approximation. After all, one might first want to find out if the tensor network is already sufficient to solve the task to the desired accuracy. If the simplest first-order method is employed, it is not clear how low the energy can go. This is particularly relevant together with the first point, where even though low entropy initializations are sub-optimal, it is not clear whether the significantly worse initial energy won't destroy the advantages. I understand that better methods can be plugged in, but this trade-off between initial energy and initial entropy is highly relevant.

These comments bring me to my major problem with the current work. After an interesting first half, it simply ends too early.

The arguments around the manifold of reachable states is rather unconvincing because there are many possible interpretations of the results around Figure 6. I don't necessarily believe the authors are wrong in their intuition about the volume of reachable space, but the results are ambiguous. I, for one, am skeptical that more reachable states alone enhance QAOA optimization, given that this has not worked well before in the context of expressivity induced barren plateaus. The solutions to these problems can trivially be prepared but are hard to find, which is why I find the intuition of reachable volume harder to believe.

Furthermore, the numerics for the proposed hybrid framework are not thorough. Figure 7 and Figure 8 appear to present individual cases at small scale and without direct comparison to other approaches. I find it confusing why the spirit of Figure 7 and 8 are not combined into one more significant comparison that nicely concludes this paper. Figure 7 also doesn't seem to follow from earlier results in this work where higher temperatures always do better. I assume this is because lower temperature reaches significantly better energies. It is not clear to me why the authors seem to avoid more meaningful numerics of the whole framework.

A few minor comments:

- Any plot or insights covering realistic energies and entropies based on maximum bond dimensions would be nice.
  - Please emphasize more throughout this work when figures were made based on \*final energies after QAOA optimization\*.
- This is often implied or weakly stated but I believe readers skimming mostly the figures will miss this part due to how much

the landscapes look like QAOA landscape slices.

- Does Figure 3 have a peak in terms of which temperature is optimal? Or is the highest entropy state at equal energy always the best?
- It is very hard for me to understand the implications or even the result of Figure 9.

## Reviewer #2

(Remarks to the Author)

Under the framework of pre-optimization of variational quantum algorithms (VQAs) using tensor networks (TNs), the authors defined a new measure, the coherence entropy, as a way to quantify the effectiveness of the initial state of the quantum approximate optimization algorithms (QAOA). The authors identified pure Gibbs states as the maximizer of coherence entropy given energy, and used matrix product state (MPS) imaginary-time evolution algorithms to generate the state. The MPS is constructed on a quantum computer through standard techniques. The pure Gibbs state is then benchmarked against other typical initializations. Finally, the authors analyze the intuition for maximizing relative coherence entropy based on the volume of reachable states, and is validated through numerics.

As far as I know, this work is novel and such coherence entropy measure for the effectiveness of initialization is not explicitly considered in QAOA by other work. The numerical benchmark is convincing, as not only several initial states including product, Hadamard, and Gaussian are compared against, the energy landscape is also studied, which convincingly shows that for states with same energy, a higher coherence entropy helps. The interpretation of the effectiveness of this measure through the principle component analysis is also reasonable. I believe this paper will definitely influence thinking in the field, especially in shaping strategies for selecting initial states in VQA algorithms and evaluating the potential benefits of classical pre-training. Additionally, coherence entropy offers an interesting perspective for exploring VQAs from a rigorous information-theoretic perspective.

I want to bring up the following minor points:

1. It may be worth mentioning that Eq.(6) only optimizes Eq.(5) if  $H$  is diagonal (classical), even though it is clear from the context of QAOA.
2. I'm confused about the relative approximation ratio. In particular, isn't Eq.(20) being **negative** implying  $E_y$  is lower than  $E_x$ , so  $y$  is better?

Overall, I recommend the acceptance of the manuscript.

## Reviewer #3

(Remarks to the Author)

A coherent approach to quantum-classical optimization

Andres Caliz et al

This paper uses a Gibbs state obtained using tensor network methods to initialise a QAOA optimisation. Varying the temperature of this initialisation partitions the effort required for the optimisation between the classical and quantum parts. Interestingly, it is demonstrated that (with appropriate scaling) a Gibbs state provides a better initialisation than a product state at the same temperature. This is argued – evidenced by use of the diagonal entropy as a measure of entanglement - to be because the Gibbs state is a superposition of a greater fraction of QAOA states.

The results – while quite heuristic – appear sound and are interesting. On balance I recommend acceptance in Communications Physics.

One minor comment that the authors might consider before publication. I suggest that sections D and E of the methods be removed. The discussion here is not really self-contained using notation and language that is only really defined in Ref[30]. It is probably not necessary for readers familiar with Ref[30] to have these sections and Ref[30] itself is more readable than Section D and E.

Version 1:

Reviewer comments:

## Reviewer #1

(Remarks to the Author)

The authors have significantly improved the clarity of their work. As before, I believe that this work offers an interesting perspective on classical pre-training in QAOA. It could be argued that the overall framework should be tested more, but I understand that it is not the author's focus. As such, I believe that this manuscript can be published in its current form.

If I could propose one final addition, it would be something akin to Figure 1b in the author's reply. This is not strictly necessary, but I found this figure very interesting. With data from actual numerical simulations, with the imaginary time and

DMRG trajectories on the Energy-Entropy surface, I believe that it could be very helpful to the readers. It would very clearly summarize the results of this work in one swoop, showcasing that DMRG yields a lower initial energy but also low entropy, and that subsequent optimization (potentially also added with arrows) from those initializations can thus be suboptimal. If the authors agree that this could enhance their work, they may consider adding such trajectories in their current Figure 2.

**Open Access** This Peer Review File 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.

In cases where reviewers are anonymous, credit should be given to 'Anonymous Referee' and the source.

The images or other third party material in this Peer Review File 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 <https://creativecommons.org/licenses/by/4.0/>

1. The one issue I have with this insight, also in Appendix A, is that it neglects the possibility to early-stop the MPS training when entropy or the bond dimension starts to drop. A comparison to such an approach would be very nice.

The reviewer is correct in pointing out that a possible approach to improve the existing methodology would be to early stop the DMRG algorithm before convergence of the algorithm, ending up with a state with higher entropy but slightly worse energy. However, we note that drops in entropy when solving the problem with the DMRG tend to be very sharp on our numerical tests, and it is unclear how to properly define the stopping points, due to oscillations in these quantities (see Figure 1). **Additionally, and as seen from the numerical results on the QAOA, irrespective of the classical algorithm used to generate the initial quantum state, the quantum part of the optimization is most “useful” (i.e. it produces the largest reduction in energy) when the entropy of the initial state is maximal.** In accordance to this result, we suggest a protocol that exploits this by approximately generating maximally coherent states classically, as otherwise the hybrid protocol is hardly justified.

That being said, we have modified the text to explicitly clarify that the quality of the optimization depends on the combined performance of each part of the hybrid protocol, which depends on the available resources. Therefore, both the classical and quantum algorithms used could be adapted accordingly (lines 548-555).




Figure 1: Behavior of the DMRG algorithm. a) It shows the evolution of the coherence entropy of the states found by the DMRG algorithm throughout the entire optimization process. b) It graphically illustrates the disparity in behavior between the DMRG algorithm and the MPO WI,II algorithm during the optimization process. While the MPO WI,II algorithm (red line) performs the optimization through imaginary time evolution, ensuring the creation of states with maximum entropy, the DMRG algorithm (brown line) carries out an optimization by navigating through the interior of the diagram, generating states that are farther from the Boltzmann boundary.

As per early-stopping strategies for the newly proposed protocol, we note that the imaginary time evolution carried out by the MPO WI,II algorithm, as indicated in Section E of Methods, involves the iterative application of the time evolution operator in its approximate form (WI or WII). When this operator is applied in the form of an MPO to the state, represented as an MPS, the state evolves by an imaginary time  $\delta\tau$  for each layer (step) of the MPO. Consequently, a smaller maximum allowed bond dimension or a smaller number of steps leads to a different imaginary evolution time, which affects the quality and temperature of the pure Gibbs states obtained. In other words, to achieve longer imaginary evolution times (associated with lower-temperature Gibbs states), both a larger bond dimension and a greater number of steps are required. Figure 9 of Appendix C presents an analysis focused on the WII version of the tensor network, in which the algorithm’s performance is evaluated for generating high-quality Gibbs states with various bond dimensions and different evolution times.

2. Though when it comes to actually applying the imaginary time evolution, the very rough approximation via the first order method is not appropriately studied. For example, it is not clear how close to the Boltzmann limit of entropy the evolution happens, or how well the solution could have been classically given this approximation.

We find the reviewer’s comment very appropriate, as indeed the article does not include any reference to the performance of the WI version of the MPO WI,II algorithm in creating Gibbs states at different temperatures. Over the course of our work, we have found that for the task of creating Gibbs states, both versions have very similar accuracy.

This is because the error introduced in the MPS truncation process, which is carried out after applying the MPO time-evolution operator, is greater than the difference in error introduced by each version of the algorithm, at least for the system dimensions used. Similarly to the results of Figure 9 of the article, if we compare the quality of the Gibbs states generated by the MPO WI,II algorithm for different values of the maximum bond dimension and different evolution times, the results are effectively identical for both the WI and WII versions. Note that such plot displays the amplitudes of all the computational basis states, given its energy, in the superposition that compose the reached approximate Gibbs state. Hence, a perfect approximation would be plotted as a straight line with a slope dependent on the temperature. The differences in the “cue” of the distributions achieved by WI and WII algorithms are non-relevant numerically.




Figure 2: Amplitude distributions illustrating the quality of the Gibbs states generated by MPO WI,II for an 8-qubit system. Figures (a) and (b) show the results obtained with the WII version of the algorithm, while (c) and (d) present the results corresponding to the WI version.

In Appendix C (lines 898-904), an additional clarification regarding the performance difference between the WI and WII versions of the algorithm has been included. Moreover, a minor modification has been made to the caption of Figure 9 to indicate that the results are almost equivalent by both versions. An explanation of the Pearson correlation coefficient has also been added to provide greater interpretability for Figure 9 (lines 887-894).

Additional text has also been included that summarizes the results of Figure 9 and interprets them in terms of the distance between the created Gibbs state and the ideal Gibbs state in the Boltzmann limit (lines 906-911).

3. This is particularly relevant together with the first point, where even though low entropy initializations are sub-optimal, it is not clear whether the significantly worse initial energy won't destroy the advantages. I understand that better methods can be plugged in, but this trade-off between initial energy and initial entropy is highly relevant.

Figure 4 has been updated to show the relative value of  $\alpha_r$  between the Gibbs states located on the boundary and the rest of the states in the entropy-energy diagram. Although the previous figure showed that Gibbs states were better than the rest, it could mistakenly suggest that all Gibbs states are equally good. Now, the new Figure 4 presents the entropy-energy diagram, where  $\alpha_r$  compares the states in the diagram with the Hamiltonian's minimum-energy state.

This new figure enhances the interpretability of the presented data. On the one hand, for a given energy value (x-axis), states with higher entropy yield better results. On the other hand, states with lower energy produce better results because they are closer to the global solution. Taken together, the new figure shows the combination of both effects, once again demonstrating that the states lying on the Boltzmann boundary are the ones that allow the quantum algorithm to maximally improve the quality of the solution. In the following figure, the heatmap presented in Figure 4 of the article is shown, where  $\alpha_r$  is used for different types of state pairs.




Figure 3: Heat map for a 6-node instance of the Max Cut problem. (a) This is the original figure included in the article, in which  $\alpha_r$  is measured for the states located in the energy-entropy diagram against the Gibbs states on the boundary. (b) For each state located within the diagram, the associated color indicates the value of the relative quality  $\alpha_r$  of the indicated state always compared against the minimum energy state  $|\psi_{E_{\min}}\rangle$ . (c) For each state in the entropy-energy diagram, a value  $\alpha_r$  is measured, which involves using both the original state and the state obtained through the QAOA optimization process. It provides a measure of the energy distance traversed by the QAOA algorithm. It is clear that as lower-energy states are used, the maximum potentially traversable distance decreases.

The heatmaps presented in Figure 3 consistently show that the states on the Boltzmann boundary maximize the performance of the QAOA algorithm in general, as they maximize the change in energy covered by the algorithm while also enabling the discovery of lower-energy solutions. Nonetheless, the sweet-spot that achieves the optimal result for any given instance and with the desired amount of quantum and classical resources will be instance, ansatz, optimizer and initialization-dependent, in general.

4. **The arguments around the manifold of reachable states is rather unconvincing because there are many possible interpretations of the results around Figure 6. I don't necessarily believe the authors are wrong in their intuition about the volume of reachable space, but the results are ambiguous. I, for one, am skeptical that more reachable states alone enhance QAOA optimization, given that this has not worked well before in the context of expressivity induced barren plateaus. The solutions to these problems can trivially be prepared but are hard to find, which is why I find the intuition of reachable volume harder to believe.**

While barren plateaus are often linked to an increasing circuit size—since deeper circuits generally introduce more parameters and can lead to an exponential vanishing of gradients—, that is not a relevant factor in our analysis, given that all our tests are performed for circuits of very similar (and small) sizes and equal number of parameters. However, it is true that barren plateaus have also been studied in the context of different initialization schemes for variational algorithms. We agree with the reviewer that we can not guarantee that the increase in expressivity is the sole factor that explains the performance improve, but further research seems to validate its relevance.

In this sense, our reasoning stems from the empirical observation that the landscape generated by a single-layer QAOA has more extreme optima (“lower low’s and higher high’s”) than one generated with low entropy states, as depicted in Figure 5 in the paper. With this in mind, and acknowledging that the existence of better optima need not imply that these optima are necessarily reached, we hypothesize this occurrence could be linked to an increased effective-volume covered, resulting in higher energetic variance of the generated states.

In order to increase the strength of our argument linked to Figure 6 of the article, the original figure has been replaced with a new one. This new figure presents the entropy-area and entropy-rank relationships of the generated manifold, and it also includes the observed variance for each of these generated manifolds. In the new figure, it can be observed that as entropy increases, the area also increases, and so does the variance in the energy landscape. We have also modified the text to include this explanation.

5. **Furthermore, the numerics for the proposed hybrid framework are not thorough. Figure 7 and Figure 8 appear to present individual cases at small scale and without direct comparison to other approaches. I find it confusing why the spirit of Figure 7 and 8 are not combined into one more significant comparison that nicely concludes this paper**

We acknowledge that the numerics on the **whole** workflow of the proposed hybrid framework are not thorough, but we believe the reviewer’s concern should be alleviated considering all the previous results shown for each of the components analysed in isolation. Each component is studied with extensive numerical tests (Figures 3,4,6,9), demonstrating that the relevance of the quantum part of the hybrid quantum-classical optimization protocol is clearly dependent on the entropy of the initial state. We make sure not to claim that the new protocol will achieve superior quality results to the previous one, but rather that the previous protocol can not take advantage of the quantum optimizer by design. In this regard, and for the size of the generated tests, the DMRG algorithm reaches the global solution of the problem when a sufficiently large bond dimension is used, and so do many other heuristic classical algorithms. On the other hand, performing larger tests requires the execution of the algorithm on real quantum hardware, which is a complex task and may require additional tweaks on the protocol depending on the capabilities of the quantum hardware used (e.g. switching from the QAOA to a Hardware Efficient Ansatz).

6. **Figure 7 also doesn't seem to follow from earlier results in this work where higher temperatures always do better. I assume this is because lower temperature reaches significantly better energies. It is not clear to me why the authors seem to avoid more meaningful numerics of the whole framework.**

Gibbs states are quantum states that have an associated temperature, given by Equation 6 of the article. Moreover, Equation 12 shows that the temperature of each Gibbs state is inversely proportional to the imaginary time evolution required to generate them. Therefore, lower-temperature Gibbs states require longer evolution times, indicating that they are the most difficult to produce.

In Figures 3.a and 3.b of the article, it is shown that lower-temperature Gibbs states yield better results. Similarly, in the new Figure 4 introduced in the article, it can be seen that lower-energy (lower-temperature) Gibbs states produce better results. This is consistent with what is presented in Figure 7, where lower-temperature states also achieve better results. Specifically, it can be observed that the state with temperature  $T = 2$  reaches higher-quality solutions than the state with  $T = 3$ , which in turn outperforms the Hadamard state, understood as a Gibbs state at infinite temperature.

**7. Any plot or insights covering realistic energies and entropies based on maximum bond dimensions would be nice.**

We attach a comparison of the entropy and energy of the states generated by the DMRG at different problem sizes and bond dimensions. As previously mentioned, the DMRG already achieves close to optimal results with very small bond dimensions ( $\chi = 4$ ). As per the entropy of the generated states, the average value remains almost 0 irrespective of the problem size. These results are specific to a particular aspect of the discussion and do not directly contribute to the broader argument of the paper. Therefore, we have opted not to include them in the final manuscript.




Figure 4: Performance analysis of the DMRG algorithm. a) It shows the coherence entropy of the solution states obtained using the DMRG algorithm for different problem sizes. b) It shows the quality of the obtained solution using the approximation ratio metric, which reflects the distance to the exact solution of the problem.

**8. Please emphasize more throughout this work when figures were made based on \*final energies after QAOA optimization\*. This is often implied of weakly stated but I believe readers skimming mostly the figures will miss this part due to how much the landscapes look like QAOA landscape slices.**

We thank the reviewer for this comment. We have modified our notation in the article such that the  $U_{\text{QAOA}}|\psi\rangle$  is now used to refer to the state generated by the QAOA algorithm after the optimization process, using the initialization state  $|\psi\rangle$ . The use of this notation has been presented (lines 378-382) and Figures 3, 4, and 5 have been updated to incorporate the new notation.

**9. Does Figure 3 have a peak in terms of which temperature is optimal? Or is the highest entropy state at equal energy always the best?**

The data obtained from the simulations show that, for states with the same energy but different entropy, the state with higher entropy tends to produce better results. Thus, there is no observed optimal temperature peak: any lower-temperature Gibbs state yields better results than a higher-temperature Gibbs

state. This statement is consistent with the findings presented in Figure 3.b and Figure 4. Additionally, we note that, in the theoretical limit of  $T \rightarrow 0$ , the Gibbs state corresponds to the ground state of the problem, and thus the global optimum of the optimization task.

**10. It is very hard for me to understand the implications or even the result of Figure 9.**

The purpose of Figure 9 is to present the performance of the MPO WI,II tensor network algorithm in constructing Gibbs states at different temperatures (associated with different evolution times). As mentioned in a previous response, for an exact Gibbs state, the plot would correspond to a straight line with a slope that depends on the temperature. The Pearson correlation coefficient is used as a quantitative metric to reinforce the qualitative observation of the plot.

**Reviewer #2**

1. **It may be worth mentioning that Eq.(6) only optimizes Eq.(5) if  $H$  is diagonal (classical), even though it is clear from the context of QAOA.**

We find the reviewer's comment very appropriate, and consequently, we have added an additional clarification in the article (lines 176–179), indicating that the constructed state solution is only valid when  $H$  is diagonal.

2. **I'm confused about the relative approximation ratio. In particular, isn't Eq.(20) being negative implying  $E_y$  is lower than  $E_x$ , so  $|y\rangle$  is better?**

Thank you very much for pointing this out, as there was a mistake in equation 20. Equation 20 and the subsequent expression (line 374) have been corrected to indicate that the absolute values of the energies are used. Now, with the corrected expression, a positive value of  $\alpha_r$  indicates a higher quality for the state  $|y\rangle$ .

**Reviewer #3**

1. **I suggest that sections D and E of the methods be removed. The discussion here is not really self-contained using notation and language that is only really defined in Ref[30].**

We believe that, although the information presented in Section D of Methods could be omitted, it is appropriate to keep it for pedagogical reasons. Moreover, Section E of Methods is not explicitly described in Ref. [30], and since a significant portion of this article focuses on Gibbs states, we consider it useful to show how they are constructed using the numerical method presented in Ref. [30]. We have revised the text to validate its self-consistency, and we included a sentence to refer the reader to the original paper for additional details (lines 241-243).