nature portfolio

Peer Review File

The phase diagram of quantum chromodynamics in one dimension on a quantum computer

Corresponding Author: Dr Norbert Linke

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)
In their manuscript “The phase diagram of quantum chromodynamics in one dimension on a quantum computer,” Than et al. employ a digital trapped-ion quantum computer that utilizes motional phonon modes as ancillary qubits to implement a hybrid sampling/variational algorithm for preparing thermal states of a lattice hypercell in SU(2)/SU(3) lattice gauge theory (LGT) in 1+1 dimensions. The authors introduce two techniques: first, the use of motional ancillae as part of the quantum circuit, and second, a projective measurement scheme tailored to the symmetries of the model.

The underlying model is the Kogut-Susskind formulation of SU(2)/SU(3) LGT in 1+1D coupled to fermionic matter. In one spatial dimension, Gauss’s law is explicitly solved, leading to a purely fermionic model with long-range interactions, which no longer is a gauge theory. Applying a Jordan-Wigner transformation maps the fermionic degrees of freedom onto qubits; however, this generally results in nonlocal N-body interactions due to the Jordan-Wigner string. To avoid these complications, the authors restrict their implementation to two-site systems, corresponding to four and six system qubits, respectively. The authors address the phase diagram of the model, i.e., how an order parameter, the chiral condensate, changes with temperature and chemical potential.

The thermal state preparation algorithm is based on a variational quantum eigensolver (VQE) framework that minimizes the free energy. The approach involves a parametrized unitary that entangles the motional ancillae with the system qubits, thereby generating a tunable probability distribution, though without entangling operations between the ancillae, and inducing a mixed state on the system. A second variational circuit is then applied to the system qubits to further optimize the state.

Finally, the authors perform measurements. The procedure whereby Gauss law is integrated out leaves the global SU(2)/SU(3) charge of the resulting fermionic model unspecified. The authors enforce that the total state lies in the singlet sector (no global charge) by applying a group-theoretical projection onto the gauge-invariant subspace.

To assess the innovation of the paper, two aspects must be considered. First is the technical achievement of realizing motional ancillae and implementing the VQE circuits. The use of motional ancillae is a particularly intriguing development, as it enhances the capabilities of trapped-ion devices by extending the accessible Hilbert space—especially important given the size limitations of ion chains. This is quite intriguing as there are phonon modes along three dimension. To enable gate operations only those along one dimension are needed, while the others can be used as additional dofs. The fact that the authors have demonstrated (some level of) their control is an important development. One can envision further developments involving multiple phonon excitations, potentially enabling the use of motional modes as an effective heat bath for the system qubits. This is novel and innovative.

Second is the choice of physics targets and the paper's contribution to advancing them. In contrast to the technical achievements, I find the paper lacking in several aspects regarding the physics it aims to address and the presentation of the model. A reader from the high-energy or nuclear physics community—the intended audience for this work—is unlikely to extract much insight from the current version. The presentation of the model, including its restrictions and physical assumptions, is insufficient and in some cases misleading. These issues require revision.

1. The discussion of the model is relegated to the Methods section and Supplementary Information, apart from a brief

abstract statement in Eq. (1). While the model itself is well-studied, it involves subtleties that are critical for correctly interpreting the results. In particular, it makes a substantial difference whether one works with the full gauge theory or with a formulation in which the Gauss laws have been integrated out, resulting in a non-gauge-invariant model. Similarly, the fermion-to-qubit mapping plays a crucial role: integrating out Gauss laws in the qubit representation typically leads to nonlocal, multi-qubit interactions. This does not mean the authors must include an exhaustive discussion of a known model, but the paper must be self-contained enough for readers to understand the essential assumptions and approximations. As it stands, crucial information is missing from the main text. For instance, it is never stated—even in the figure captions presenting the main results—that the simulations are performed at N=2 sites, which eliminates Jordan-Wigner-induced long-range interactions. This is a key simplification that needs to be clearly communicated.

2. The presentation of the projection method, and especially in Fig. 1b and 1d, is misleading. The authors integrate out Gauss laws. At this point, the model no longer has a gauge symmetry. There is a remaining global SU(2)/SU(3) charge corresponding to setting the boundary conditions of the model. However, the authors describe a projection method whereby they make “gauge-invariant measurements” and “project onto the physical space” from the “full Hilbert space”. Both statements are incorrect. The gauge symmetry is integrated out. All states in the fermionic formulation are physical. The projection is on a global charge-neutral sector. This distinction is not just semantic—it changes the physical interpretation of the results.

3. Given that the projection approach is one of the two main innovations claimed by the paper, a more detailed and accurate explanation is warranted. For example, a projector onto the singlet sector can be formally expressed via a Haar integral over group transformations (see, e.g., Kasper et al., Phys. Rev. D 107, 014506 (2023)). In contrast, Eq. (S5) in the Supplementary Information shows a projector that is diagonal in the computational basis. This suggests that one can simply post-select measurement outcomes based on basis-states, which is surprising. If I can simply postselect, why introduce a projection operator in the first place? Naively, I would have guessed that such a procedure would only be valid for Abelian symmetries. In the case of SU(2), for instance, projecting onto a state with <Sz>=0 does not preclude the global SU(2) spin being in the transverse (x–y) plane. Is there an additional assumption or structure in the model that justifies this basis-diagonal projection? If so, it should be clearly stated.

A lesser, but still relevant point:

a. The primary motivation is to overcome the sign problem of Monte Carlo methods in Lattice QCD at finite density, which is crucial for mapping out the QCD phase diagram. While the authors discuss classical tensor network techniques, they overlook some of the extensive existing literature on this longstanding problem. Given that the target audience includes researchers in this well-established field, it is important to cite relevant prior approaches—e.g., complexification and path deformation methods, reweighting techniques, and Taylor expansion methods—or a few review articles, from that field or more general.

On balance, I am torn between recognizing the clear innovation introduced by the use of motional ancillae and the implementation of the VQE algorithm, which opens up several promising avenues for future exploration, versus discussion of the physics and its presentation. The control of motional modes, even if not yet universal, is an important technical achievement that could plausibly justify publication in Nature Communications. Regarding the physics target, it feels almost as if, had the authors focused on a more “standard” model, such as the transverse field Ising model, the core message of the paper might have been more effectively conveyed. That said, the chosen target—SU(2)/SU(3) in one dimension—is certainly a compelling and important one, and the title, “The phase diagram of quantum chromodynamics in one dimension on a quantum computer,” reflects this ambition.

As it stands, unless significant revisions are made, I am unable to recommend the manuscript for publication.

Reviewer #2

(Remarks to the Author)
I have read this paper with interest and find the theme, conclusions and objectives of the research very interesting. It addresses a difficult and important problem, namely how to best carry out numerical/experimental simulations of the physics of QCD. The authors propose a numerical approach, which is based for the moment on studying lower dimensional versions of QCD, that is theories in 1D which nonetheless possess SU(2) and SU(3) symmetries. The numerical method that is proposed is inspired by quantum computation protocols and transfers nicely to cold atom experiments. My understanding is that this suggests a pathway by which the physics of QCD may be studied by carrying out measurements on cold atom systems, which is a very interesting viewpoint.

My view is that the results of the work are important and worthy of publication in Nature Communications. The paper is nicely written, the results are well explained and the measurement and preparation protocol presented in a lot of detail, through useful diagrams, detail descriptions in the body of the paper and more detailed appendices.

There is one aspect of the work which I think can be improved and would be helpful to future readers like myself who are not experts on numerical approaches to simulating quantum systems but are at least familiar with some of them. If I understand correctly the explanations and diagrams, the aim of the protocol is to have access to the expectation values of local observables in thermal states. There are many different methods out there, especially for 1D systems, to compute such expectation values, from tDMRG to Matrix Product States or even a hydrodynamic approach, not to mention methods that are specific to systems with additional symmetries, such as in the presence of integrability. In essence, what I miss here is a

more detailed discussion of how this methodology is superior to others, or of how it can uniquely address the problem of interest. This is not clear enough for me from reading the paper, but I think should be made clearer given the claims the authors make.

Reviewer #3

(Remarks to the Author)
The paper "The phase diagram of quantum chromodynamics in one dimension on a quantum computer" presents a pioneering experimental study of thermal states in lattice gauge theories (LGTs), using a trapped-ion quantum computer. The authors target the quantum chromodynamics (QCD) phase diagram in one dimension by simulating SU(2) and SU(3) gauge theories at finite temperature and chemical potential. The work addresses two core challenges in simulating thermal states of LGTs: (1) preparing mixed thermal states, and (2) ensuring gauge invariance. They overcome these by introducing motional ancilla to encode thermal probabilities and by performing gauge-invariant measurements that evaluate the ratio of two observables instead of explicitly enforcing Gauss's law in the quantum circuit.

Research that develops methods for quantum simulation, particularly in less-explored areas like lattice gauge theory, using quantum systems is highly meaningful. In general, I recommend the publication of the paper to Nature Communications. However, I feel many basic parts of the quantum simulation for the LGT are not well explained as listed below. I hope the authors address these points properly before the publication.

1. The physical interpretation of the Hamiltonian in this work is not entirely clear. While lattice gauge theories (LGTs) are known to capture a variety of physical phenomena, it remains unclear which specific phenomenon is being addressed in this simulation. In particular, the meaning of the resulting phase diagram should be clarified—what exactly does it represent, and how is it connected to known or expected phases in QCD or other gauge theories? The choice of the observable—the chiral condensate—as the order parameter also requires further explanation. Given a chemical potential, what is the explicit form of the ground state? The terminology of "symmetry-broken" and "symmetry-restored" phases is not self-evident. For instance, in Fig. 3d, the vacuum state appears to be the ground state at low chemical potential, and the baryon state dominates at high chemical potential, but it is not obvious why this is the case. Clarifying these points would significantly improve the clarity and impact of the results.

2. The main point of this work lies in the efficient preparation of thermal states using motional mode ancilla, but a more detailed explanation of this approach is necessary. While the use of system–motional ancilla interactions to evaluate the entropy, as shown in Eq. (4), without direct measurement of the motional modes is well explained, it seems to be a separate issue from the actual preparation of the thermal state in Eq. (2). In particular, it is not clearly described how thermal states at specific temperatures are prepared. Since the system qubit states are generally not eigenstates of the Hamiltonian in Eq. (1), it is unclear whether the proposed method using motional modes is capable of preparing the desired thermal state at a given temperature.

3. It would be helpful if the authors provided some discussion on the importance of preparing thermal states in quantum simulators. In particular, it would be valuable to explain the challenges of incorporating finite-temperature effects in classical computations compared to zero-temperature calculations, and how quantum simulation may offer advantages in this regard. A comparison of the additional computational cost typically required, especially how this cost scales with system size beyond a single unit cell, in both classical and quantum approaches, would greatly enhance the context and significance of the present work. Additionally, it would strengthen the impact of the paper if the authors discussed whether their thermal state preparation method using motional ancillae is scalable to larger systems.

4. Regarding the measurement method, it would be helpful if the authors discussed the general applicability of the approach presented in Eq. (5). For the specific case of the chosen order parameter—the chiral condensate in SU(2) and SU(3)—the expectation value can be evaluated using only qubit measurements in the computational basis. However, it remains unclear how broadly this technique can be applied. For instance, can this projection-based measurement scheme be extended to systems with more lattice sites beyond a single unit cell? Is it applicable in higher-dimensional settings such as 2D or 3D lattice gauge theories? It would also be valuable to clarify to what extent this approach can serve as a substitute for explicitly implementing Gauss's law in the system.

5. When varying the chemical potential μ, what changes are required in the quantum simulation? Is μ simply treated as a parameter in the classical optimization procedure for minimizing the free energy, or does it necessitate adjustments in the quantum circuit as well?

6. In Fig. 3, it appears that only a single experimental data point is presented at $\mu \approx 2.0$. Could the authors clarify why only this point was chosen? What are the main challenges in obtaining experimental results for other values of μ, and how do they differ from the SU(2) case where multiple data points were provided?

7. The variational circuit U_S shown in Fig. 1c seems relatively complex and resource-intensive. Beyond a brief mention of noisy numerical simulations, there is little discussion of the required gate fidelities and the actual performance of the quantum gates used in the experiment. How good were the quantum operations? What are the required fidelities for the method to remain accurate? Were any error mitigation techniques employed to improve performance? A more detailed discussion of these aspects would be highly valuable.

Version 1:

Reviewer comments:

Reviewer #1

(Remarks to the Author)
The authors have satisfactorily addressed my previous comments and significantly improved the presentation, especially of the "physics application" component. With these changes, and given the innovation in the technical achievement of realizing motional ancillae, this is sufficient to justify publication in Nature Communications.

I have one remaining recommendation. Although the authors have clarified the role of gauge symmetry in their model, they continue to use the term 'gauge-singlet', where 'charge-singlet' would be more appropriate. (Strictly speaking, in this work, the projection is applied to a global, not a local symmetry.) I don't mean to nitpick, but let me frame the concern constructively: consider the Heisenberg model, which has a global non-Abelian SU(2) symmetry. While this is not a gauge symmetry at all, the projection technique proposed by the authors might still be applicable. Referring to the relevant states as gauge-singlets rather than charge-singlets could unintentionally suggest a narrower scope of applicability than actually exists.

Reviewer #2

(Remarks to the Author)
I have reviewed the second version of this paper, together with the authors detailed answers to all three sets of referee comments. My original report was the most positive of the three and my original decision was to accept the paper subject to minor improvements. The authors have responded adequately to my questions and made small changes in the paper as a result.

The other two referees are clearly more experienced in numerical methods and had many comments and criticisms of the method and models studied. Looking at the authors' answers and changes to the paper, I feel actually more confident than before in my decision to accept the paper. The main weakness of my report was indeed that I am no expert on numerical methods. As such my judgment was based on an appreciation of the importance of the problem that is being dealt with rather than a deep understanding of the methodology proposed. I am therefore very pleased that the other referees have been able to ask deeper questions about the methodology.

My feeling is that the authors have dealt well with those questions and strengthen their paper as a result. Therefore I stand by my original decision to accept the paper for publication.

Reviewer #3

(Remarks to the Author)
Overall, the authors have responded well, and I have no further objections to the publication of the paper. However, I would appreciate it if the following points could also be clarified:

Regarding thermal state preparation, the current explanation still leaves some confusion. It gives the impression that one would need to know all eigenstates in advance, which would essentially mean the problem is already solved.

In the fidelity figures, the case of SU(3) appears to show more serious deviations—would it be possible to provide additional data or discussion on this point? Also, since larger discrepancies occur in the low chemical potential region, I wonder if the experiment could have focused more on that regime.

We sincerely thank the editor and all three reviewers for their time, thoughtful evaluation, and constructive feedback on our manuscript. Below, we reproduce the reviewers' comments and provide our detailed responses to each point. All changes made in the text are denoted in blue color.

.....

Reviewer 1: In their manuscript "The phase diagram of quantum chromodynamics in one dimension on a quantum computer," Than et al. employ a digital trapped-ion quantum computer that utilises motional phonon modes as ancillary qubits to implement a hybrid sampling/variational algorithm for preparing thermal states of a lattice hypercell in SU(2)/SU(3) lattice gauge theory (LGT) in 1+1 dimensions. The authors introduce two techniques: first, the use of motional ancillae as part of the quantum circuit, and second, a projective measurement scheme tailored to the symmetries of the model.

The underlying model is the Kogut-Susskind formulation of SU(2)/SU(3) LGT in 1+1D coupled to fermionic matter. In one spatial dimension, Gauss's law is explicitly solved, leading to a purely fermionic model with long-range interactions, which no longer is a gauge theory. Applying a Jordan-Wigner transformation maps the fermionic degrees of freedom onto qubits; however, this generally results in nonlocal N-body interactions due to the Jordan-Wigner string. To avoid these complications, the authors restrict their implementation to two-site systems, corresponding to four and six system qubits, respectively.

We appreciate this observation of the reviewer. In the main text (second paragraph of Page 3) we have now introduced a comment about this (also explained below in the response to comment 1 of the reviewer) for clarity about the scope of this article:

- We state now more clearly that long-range interactions appear in the Hamiltonian as a consequence of the Jordan-Wigner (JW) transformation, and we explain that the many-bodyness and type of this interaction does not change with system size, i.e., we have four-body nonlocal interaction for SU(2) and six-body non-local interaction for SU(3). These long-range interactions disappear for a single unit cell.
- We also added a statement explaining that the resulting long-range interactions can be implemented natively on a trapped-ion device (such as the one used for the experiment in this article) due to the inherent all-to-all connectivity among all the ions in the trap.

Since the phase diagram remains qualitatively the same, in this context studying a unit cell provides us with necessary physical insights.

Reviewer 1: The authors address the phase diagram of the model, i.e., how an order parameter, the chiral condensate, changes with temperature and chemical potential. The thermal state preparation algorithm is based on a variational quantum eigensolver (VQE) framework that minimizes the free energy. The approach involves a parametrized unitary that entangles the motional ancillae with the system qubits, thereby generating a tunable probability distribution, though without entangling operations between the ancillae, and inducing a mixed state on the system. A second variational circuit is then applied to the system qubits to further optimise the state.

Finally, the authors perform measurements. The procedure whereby Gauss law is integrated out leaves the global SU(2)/SU(3) charge of the resulting fermionic model unspecified. The authors enforce that the total state lies in the singlet sector (no global charge) by applying a group-theoretical projection onto the gauge-invariant subspace.

To assess the innovation of the paper, two aspects must be considered. First is the technical achievement of realising motional ancillae and implementing the VQE circuits. The use of motional ancillae is a particularly intriguing development, as it enhances the capabilities of

trapped-ion devices by extending the accessible Hilbert space—especially important given the size limitations of ion chains. This is quite intriguing as there are phonon modes along three dimension. To enable gate operations only those along one dimension are needed, while the others can be used as additional dofs. The fact that the authors have demonstrated (some level of) their control is an important development. One can envision further developments involving multiple phonon excitations, potentially enabling the use of motional modes as an effective heat bath for the system qubits. This is novel and innovative.

We thank the reviewer for this positive and appreciative remark.

Reviewer 1: Second is the choice of physics targets and the paper's contribution to advancing them. In contrast to the technical achievements, I find the paper lacking in several aspects regarding the physics it aims to address and the presentation of the model. A reader from the high-energy or nuclear physics community—the intended audience for this work—is unlikely to extract much insight from the current version. The presentation of the model, including its restrictions and physical assumptions, is insufficient and in some cases misleading. These issues require revision.

We thank the reviewer for the detailed and stimulating feedback. We have addressed each of the comments below and revised the manuscript accordingly.

Reviewer 1: The discussion of the model is relegated to the Methods section and Supplementary Information, apart from a brief abstract statement in Eq. (1). While the model itself is well-studied, it involves subtleties that are critical for correctly interpreting the results. In particular, it makes a substantial difference whether one works with the full gauge theory or with a formulation in which the Gauss laws have been integrated out, resulting in a non-gauge-invariant model. Similarly, the fermion-to-qubit mapping plays a crucial role: integrating out Gauss laws in the qubit representation typically leads to nonlocal, multi-qubit interactions. This does not mean the authors must include an exhaustive discussion of a known model, but the paper must be self-contained enough for readers to understand the essential assumptions and approximations. As it stands, crucial information is missing from the main text. For instance, it is never stated—even in the figure captions presenting the main results—that the simulations are performed at N=2 sites, which eliminates Jordan-Wigner-induced long-range interactions. This is a key simplification that needs to be clearly communicated.

Our reply: We thank the reviewer for pointing out the gap in presentation and for the suggestions to make the article more accessible and self-contained. We have now introduced the following key information in the main text (second paragraph of page 3):

- A consequence of the JW transformation is the appearance of long range four-body interactions for SU(2) and six-body interactions for SU(3). In this work we have considered a single unit cell for both SU(2) and SU(3), where these long-range interactions are absent.
- We have added a sentence stating that the many-bodyness of the interactions does not scale with the system size and implementing the long-range interactions is not a roadblock for trapped ion devices due to the intrinsic all-to-all connection between the ions in the trap.
- We have also added a clarification about the unit cell in the figure captions for both SU(2) and SU(3) phase diagrams.

In addition, we would like to mention that the phase diagram does not exhibit any qualitative change with the increase in system size and both the projection technique and the ansatz circuit

can be straightforwardly scaled to include more than one unit cell.

As per the suggestions of the reviewer, we have also added more details about the SU(2) and SU(3) models in the Supplementary Information (SI) S.III and S.IV for the sake of completeness and accessibility. We now have the following details included in the SI:

- The explicit form of the Hamiltonian in terms of fermionic and gauge-fields. Gauge-transformation to eliminate local gauge degrees of freedom.
- JW transformation to map the fermionic Hamiltonian to a qubit Hamiltonian.
- In the text following Eqs. (S17) and (S32), we have also indicated the long range interactions that appear due to the JW transformations and mentioned the absence of them in the case of unit cells.

> **Reviewer 1:** The presentation of the projection method, and especially in Fig. 1b and 1d, is misleading. The authors integrate out Gauss laws. At this point, the model no longer has a gauge symmetry. There is a remaining global SU(2)/SU(3) charge corresponding to setting the boundary conditions of the model. However, the authors describe a projection method whereby they make "gauge-invariant measurements" and "project onto the physical space" from the "full Hilbert space". Both statements are incorrect. The gauge symmetry is integrated out. All states in the fermionic formulation are physical. The projection is on a global charge-neutral sector. This distinction is not just semantic—it changes the physical interpretation of the results.

Our reply: We thank the reviewer for this insightful comment and agree with the reviewer about the usage of the specific phrases. The gauge-singlet or color-neutrality condition indeed comes from the extra assumption that there is no background charge, i.e., the boundary condition assumes that the global charge is zero. Originally, we defined the states which satisfy this boundary condition as 'physical states'. The projection operator ensures that we measure the expectation value on the gauge-singlet subspace. We do recognise that the use of the terms 'physical space' and 'gauge-invariant measurements' are ambiguous. So, to avoid any confusion, we have replaced them with gauge-singlet/color-neutral subspace and gauge-singlet measurements throughout the text and figures. We have also clarified this distinction throughout the whole article and made changes wherever necessary. All the changes made in the article regarding this particular issue is highlighted with bold fonts as there are numerous changes associated with this comment. We have also replaced the notations with subscript *phys* to notations with subscript 0 denoting singlet subspace.

Ensuring the gauge-singlet condition for non-Abelian gauge theories is challenging and is usually implemented by adding penalty terms to the cost function or using a symmetry-preserving circuit. We offer an alternative approach, which depends on the projection method, where the effects of the gauge-singlet constraints are implemented at the stage of observable measurement. This key idea remains valid even after the changes mentioned in the paragraph above, and all our key physical insights about the phase diagram remain unaffected. The expectation value of the chiral condensate is obtained on the gauge-singlet subspace, which was our original goal.

In the case of one spatial dimension we have used the projection method after eliminating all the gauge degrees of freedom, which makes the projection operator depend only on the global charge operators $\hat{Q}^i_{tot}$. For systems with more than one spatial dimension, the projection method can be generalised to enforce Gauss laws on each vertex as well as ensure global charge constraints. The projection operator defined in such a way is also implemented on the measurement stage and avoids explicit implementation of the Gauss laws. In that sense, the projection technique can be used to implement gauge-invariance in conjunction with the gauge-singlet condition. However, in the context and scope of this paper, the projection technique is limited to only enforcing the global charge constraint and we have made the necessary changes suggested by the reviewer.

Reviewer 1: Given that the projection approach is one of the two main innovations claimed by the paper, a more detailed and accurate explanation is warranted. For example, a projector onto the singlet sector can be formally expressed via a Haar integral over group transformations (see, e.g., Kasper et al., Phys. Rev. D 107, 014506 (2023)). In contrast, Eq. (S5) in the Supplementary Information shows a projector that is diagonal in the computational basis. This suggests that one can simply post-select measurement outcomes based on basis-states, which is surprising. If I can simply postselect, why introduce a projection operator in the first place? Naively, I would have guessed that such a procedure would only be valid for Abelian symmetries. In the case of SU(2), for instance, projecting onto a state with $\langle S_z \rangle = 0$ does not preclude the global SU(2) spin being in the transverse $(x-y)$ plane. Is there an additional assumption or structure in the model that justifies this basis-diagonal projection? If so, it should be clearly stated.

Our reply: We thank the reviewer for this very important and conceptual question. The goal of the projection technique we used here is to obtain observable expectation value as if it was obtained by performing measurements on a gauge-singlet density matrix. We emphasise that we are not projecting the density matrix on the gauge-singlet subspace explicitly during the experiment, that is at no point of the protocol, we have access to the gauge-singlet density matrix. Rather we want to ensure that we get the correct expectation value of an observable $\hat{O}$ despite having a density matrix which is not gauge-singlet.

If we had a way to prepare the gauge-singlet density matrix $\hat{\rho}_0$, and we wanted to measure the chiral condensate, which is a diagonal observable, we would need to perform z-basis measurements on the gauge-singlet density matrix. However, in the experiment, we do not have access to $\hat{\rho}_0$. In our protocol, since both the projector $\hat{K}$ and $\hat{O}$ are diagonal, we only need to perform z-basis measurements for determining expectation values of $\hat{O}\hat{K}$ and $\hat{K}$ on the density matrix we prepared on the device (which is not gauge-singlet). A ratio of these two quantities give us the correct expectation value $\mathrm{Tr}(\hat{\rho}_0\hat{O})$. In other words, we are not performing post-selection on the basis of diagonal measurements to prepare the gauge-singlet density matrix, but the measurements are performed to obtain expectation values of $\hat{O}\hat{K}$ and $\hat{K}$. If we had selected any observable $\hat{O}$, which is not diagonal in the computational basis, $\hat{O}\hat{K}$ would also be non-diagonal. This would require us to perform non-diagonal measurements on the system to recover $\mathrm{Tr}(\hat{\rho}_0\hat{O})$. Since this was not clear before, we have added a sentence in the main text clarifying this (last sentence of the paragraph following Eq. (5) on page 4).

The diagonal form of the projector, as shown in Eq. (S5), arises from our specific focus on obtaining the correct expectation values of observables rather than obtaining the gauge-singlet density matrix. In contrast, the projection operator defined in Eq. (S1), which projects the density matrix onto the gauge-singlet subspace, is generally not diagonal in the computational basis. This supports the reviewer's intuition that a purely diagonal operator would not suffice to construct a gauge-singlet density matrix for a non-Abelian gauge group. However, when computing expectation values, the projection operator appears within a trace. In this context, it is possible to diagonalize the group element $\hat{U}$ in Eq. (S1) via a similarity transformation (see Ref. [S3] for more details), which leaves the trace invariant. Consequently, for the purposes of observable measurements, it is sufficient to parameterize the Cartan subgroup rather than the full SU(N) group when defining the projection operator. This parametrization yields a diagonal form of the projector, which appears within the trace operation. We have expanded the existing paragraph after Eq. (S1) to clarify this point.

In Kasper et al., Phys. Rev. D 107, 014506 (2023), the authors use a group averaging method for a given observable, which is different from the projection operator based method considered here. In our method, we do not have a modified averaged operator, rather our observable is measured as the ratio of two observables, which is based on the projection to irreducible representations from a reducible representation, a known technique in group representation theory.

> **Reviewer 1:** A lesser, but still relevant point:
>
> The primary motivation is to overcome the sign problem of Monte Carlo methods in Lattice QCD at finite density, which is crucial for mapping out the QCD phase diagram. While the authors discuss classical tensor network techniques, they overlook some of the extensive existing literature on this longstanding problem. Given that the target audience includes researchers in this well-established field, it is important to cite relevant prior approaches—e.g., complexification and path deformation methods, reweighting techniques, and Taylor expansion methods—or a few review articles, from that field or more general.

Our reply: We thank the reviewer for bringing this to our attention. We have now added several review articles as citations at the end of the existing sentence 'Several strategies... QCD phase diagram' (first sentence of the second paragraph in the introduction). The review articles cover extensive references for the approaches mentioned by the reviewer in the comment and the current status of the Lagrangian formulation of lattice QCD.

> **Reviewer 1:** On balance, I am torn between recognizing the clear innovation introduced by the use of motional ancillae and the implementation of the VQE algorithm, which opens up several promising avenues for future exploration, versus discussion of the physics and its presentation. The control of motional modes, even if not yet universal, is an important technical achievement that could plausibly justify publication in Nature Communications. Regarding the physics target, it feels almost as if, had the authors focused on a more "standard" model, such as the transverse field Ising model, the core message of the paper might have been more effectively conveyed.

Our reply: We agree that using motional modes as an ancilla register can be extended to more general problems. Preparing thermal states for an Ising model using the motional modes would indeed demonstrate the utility of this method. However, we took the opportunity to instead take a first step towards solving one of the key open problems in physics: phase diagram of QCD. This is not only an important problem but has the potential to show the advantage of using quantum computers for lattice gauge theory simulations. As a fundamental building block, our experiment captures the nature of the phase diagram of a (1+1)-D QCD unit cell, which the reviewer agrees below as a compelling target. This target, instead of the Ising model, also required the introduction of the gauge-singlet measurement technique, adapted from group representation theory, which has broad potential applicability for LGT simulations (both classical and quantum). Our article thus advances the frontier in two directions: general thermal state preparations on a quantum computer and simulation of non-Abelian lattice gauge theories.

> **Reviewer 1:** That said, the chosen target—SU(2)/SU(3) in one dimension—is certainly a compelling and important one, and the title, "The phase diagram of quantum chromodynamics in one dimension on a quantum computer," reflects this ambition.

We thank the reviewer for recognising the significance of the problem addressed in this work and for offering detailed and constructive feedback to improve its clarity. We have revised the manuscript accordingly to address all the points raised, and we believe that the updated version meets the standards for publication in Nature Communications.

.....

> **Reviewer 2:** I have read this paper with interest and find the theme, conclusions and objectives of the research very interesting. It addresses a difficult and important problem, namely how to best carry out numerical/experimental simulations of the physics of QCD. The authors propose a numerical approach, which is based for the moment on studying lower dimensional versions of QCD, that is theories in 1D which nonetheless possess SU(2) and SU(3) symmetries. The numerical method that is proposed is inspired by quantum computation protocols and transfers nicely to cold atom experiments. My understanding is that this suggests a pathway by which the physics of QCD may be studied by carrying out measurements on cold atom systems, which is a very interesting viewpoint.
>
> My view is that the results of the work are important and worthy of publication in Nature Communications. The paper is nicely written, the results are well explained and the measurement and preparation protocol presented in a lot of detail, through useful diagrams, detail descriptions in the body of the paper and more detailed appendices.

We thank the reviewer for the very positive comments on our work.

> **Reviewer 2:** There is one aspect of the work which I think can be improved and would be helpful to future readers like myself who are not experts on numerical approaches to simulating quantum systems but are at least familiar with some of them. If I understand correctly the explanations and diagrams, the aim of the protocol is to have access to the expectation values of local observables in thermal states.
>
> There are many different methods out there, especially for 1D systems, to compute such expectation values, from tDMRG to Matrix Product States or even a hydrodynamic approach, not to mention methods that are specific to systems with additional symmetries, such as in the presence of integrability. In essence, what I miss here is a more detailed discussion of how this methodology is superior to others, or of how it can uniquely address the problem of interest. This is not clear enough for me from reading the paper, but I think should be made clearer given the claims the authors make.

Our reply: We thank the reviewer for the remarks and the question. We would like to emphasise that, in contrast to existing Lagrangian-based approaches, quantum simulation does not suffer from the sign problem when applied to lattice QCD at finite chemical potential.

As the reviewer noted, tensor network methods—also formulated within the Hamiltonian framework—have been successfully used to simulate lattice gauge theories. We acknowledge this in the introduction and have now cited relevant works accordingly (Refs. [9-11] in the Introduction). In this work, we present an alternative Hamiltonian-based approach that leverages quantum computers. One key advantage of quantum computing in this context is that extending protocols to higher dimensions does not pose a fundamental conceptual barrier; rather, the challenge comes from resources required for the quantum device and mitigating noise. In contrast, generalising tensor network methods to higher spatial dimensions is more challenging due to the exponential growth of the bond dimension. We have added a comment in the fourth paragraph of the Discussion section of the main text regarding this.

That said, both approaches are still in an exploratory phase, and active research is ongoing to address the limitations inherent in each. As such, it is currently difficult to predict which method will ultimately prove more practical or advantageous. Our intent is to offer our protocol as a complementary alternative, rather than to claim superiority over existing Hamiltonian-based techniques.

Furthermore, we note that the projection method introduced in our numerical and experimental simulations is not specific to quantum computing—it is also compatible with classical simulation techniques such as tensor networks. On the experimental side, we carried out the variational optimisation on a physical trapped-ion quantum device, using the motional modes of the trapped

ions as a qubit register. This approach not only enables lattice QCD demonstrations but also opens the door to broader applications in quantum many-body physics. The use of motional modes for quantum information processing directly highlights a promising direction for trapped-ion platforms, providing increased computational power without requiring an increase in the physical system size.

In this context, we also refer to our response to Comment 3 by Reviewer 3, which addresses a closely related point.

We thank the reviewer again for taking the time to assess our manuscript and providing constructive feedback to improve our presentation. We believe the revised version addresses all the concerns raised and further supports the reviewer's recommendation for the publication of our work.

.....

> **Reviewer 3:** The paper "The phase diagram of quantum chromodynamics in one dimension on a quantum computer" presents a pioneering experimental study of thermal states in lattice gauge theories (LGTs), using a trapped-ion quantum computer. The authors target the quantum chromodynamics (QCD) phase diagram in one dimension by simulating SU(2) and SU(3) gauge theories at finite temperature and chemical potential. The work addresses two core challenges in simulating thermal states of LGTs: (1) preparing mixed thermal states, and (2) ensuring gauge invariance. They overcome these by introducing motional ancilla to encode thermal probabilities and by performing gauge-invariant measurements that evaluate the ratio of two observables instead of explicitly enforcing Gauss's law in the quantum circuit.
>
> Research that develops methods for quantum simulation, particularly in less-explored areas like lattice gauge theory, using quantum systems is highly meaningful. In general, I recommend the publication of the paper to Nature Communications.

We thank the reviewer for the kind appreciation of our work and the positive assessment of the manuscript.

> **Reviewer 3:** However, I feel many basic parts of the quantum simulation for the LGT are not well explained as listed below. I hope the authors address these points properly before the publication.

We thank the reviewer for the detailed and constructive feedback. Below, we provide a comprehensive response to each of the comments and have revised the manuscript accordingly.

> **Reviewer 3:** The physical interpretation of the Hamiltonian in this work is not entirely clear. While lattice gauge theories (LGTs) are known to capture a variety of physical phenomena, it remains unclear which specific phenomenon is being addressed in this simulation. In particular, the meaning of the resulting phase diagram should be clarified—what exactly does it represent, and how is it connected to known or expected phases in QCD or other gauge theories? The choice of the observable—the chiral condensate—as the order parameter also requires further explanation. Given a chemical potential, what is the explicit form of the ground state? The terminology of "symmetry-broken" and "symmetry-restored" phases is not self-evident. For instance, in Fig. 3d, the vacuum state appears to be the ground state at low chemical potential, and the baryon state dominates at high chemical potential, but it is not obvious why this is the case. Clarifying these points would significantly improve the clarity and impact of the results.

Our reply: We thank the reviewer for the suggestions to improve the significance and clarity of the results in the paper. The reviewer's comments can be broken down to three main questions, which we explain below.

Q1. What does the phase diagram mean and how is it connected to known or expected phases in QCD? In the same context, what is the physical significance of the chiral condensate?

Our reply: In our experiment, we choose the chiral condensate as an order parameter for probing the phase diagram in (1+1)-D. In one spatial dimension, the chiral condensate is a bilinear fermionic operator proportional to the mass term and its explicit form is given in Eq. (16) for SU(2) and in Eq. (25) for SU(3). The baryon state (see Fig. 2e and 3e), which is favored at high chemical potential, yields a zero value for the chiral condensate operator, whereas the physical vacuum state, which is energetically favored at low chemical potential leads to a large negative chiral condensate value. This transition from negative to zero chiral condensate value as the chemical potential is increased, is reminiscent of the transition from the chiral symmetry breaking phase to the chiral symmetry restored phase for the (3+1)-D QCD phase diagram. Due to the sign problem, the precise nature and smoothness of the transition in (3+1)-dimensional QCD remain unresolved, leaving open the possibility of discovering novel phases of matter at extreme values of matter densities (as in neutron stars or the early universe). Motivated by this fundamental open question, our experiment aims to capture the corresponding transition in the chiral condensate for (1+1)-D QCD.

We would also like to note that the chiral condensate has previously been used as an order parameter in studies of the phase diagram of (1+1)-dimensional Abelian lattice gauge theories—for example, $\mathbb{Z}_2$ in Physical Review Letters 131, 081901 (2023) and $U(1)$ in Physical Review D 93, 094512 (2016). This choice is a standard practice within the lattice gauge theory community.

In the main text, we have replaced the last paragraph of the section 'Thermal states in gauge theories' (on page 3) to clarify the aim of the experiment and the motivation behind choosing chiral condensate as the order parameter.

Q2. How does the chemical potential determine the density matrix?

Our reply: The chemical potential governs the statistical weights with which energy eigenstates contribute to the density matrix in the grand-canonical ensemble. The probability of an eigenstate with energy $\tilde{E}_n$ at zero chemical potential and baryon number B (expectation value of the operator $\hat{H}_{chem}$) appearing in the thermal mixture is given by the Boltzmann factor $e^{-\beta(\tilde{E}_n - \mu B)}$. A large value of μ favors states with higher baryon number. In the context of a single unit cell studied in the experiment, the maximum baryon number corresponds to a state where all particle sites are occupied, i.e., a single baryon state. Consequently, at high chemical potential, the density matrix is dominated by this baryon state. As the chemical potential is lowered, eigenstates with lower Baryon number start to contribute more. At sufficiently low μ, the energy term dominates the Boltzmann factor, and the physical vacuum state becomes the most significant contributor to the thermal mixture.

We have expanded the discussion on the main text (first paragraph under the quasi-heading 'SU(2) thermal states' on page 6) to include more details about the physical significance of the negative and zero chiral condensate values and the corresponding phases.

Q3. What does symmetry-broken and symmetry-restored phase mean?

Our reply: As mentioned in the answer of Q1 of this comment and justified in the answer of Q2, the value of chiral condensate at high chemical potential is zero, which belongs to a baryon-dominated density matrix. At a very low chemical potential value, a vacuum-dominated density matrix becomes energetically more favorable, which yields a negative chiral condensate value. This transition from negative to zero chiral condensate value as chemical potential is increased is reminiscent of the QCD chiral symmetry breaking at low chemical potential and chiral symmetry restoration at high chemical potential. This is why we call these two phases symmetry-broken and symmetry-restored phase. However, we do recognise that this might be a bit misleading, and avoid using these exact words. Instead, we say the transition happens from a vacuum-dominated phase to a baryon-dominated phase.

We have expanded the paragraph at the end of the section 'Thermal states in gauge theories' to incorporate our comment.

> **Reviewer 3:** The main point of this work lies in the efficient preparation of thermal states using motional mode ancilla, but a more detailed explanation of this approach is necessary. While the use of system–motional ancilla interactions to evaluate the entropy, as shown in Eq. (4), without direct measurement of the motional modes is well explained, it seems to be a separate issue from the actual preparation of the thermal state in Eq. (2). In particular, it is not clearly described how thermal states at specific temperatures are prepared. Since the system qubit states are generally not eigenstates of the Hamiltonian in Eq. (1), it is unclear whether the proposed method using motional modes is capable of preparing the desired thermal state at a given temperature.

Our reply: We thank the reviewer for this comment. For a given temperature, the Gibbs state in the full Hilbert space can be decomposed as $\hat{\rho}_G = \sum_n p_n \left|\psi_n\right\rangle \left\langle\psi_n\right|$. In our protocol, the motional ancillae are used to prepare the probability distribution p_n with which each bitstring $|n\rangle$ is sampled. The purpose of the CNOT gates are to ensure that a bitstring $|n\rangle$ is sampled with a parametrized probability and the system circuit is initialized with this bitstring $|n\rangle$. The circuit on the system register then aims to prepare the state $|\psi_n\rangle$. As the free energy is minimised, an expressive variational ansatz will asymptotically reach the Gibbs state. This means if the VQE is successful, in the end, we will have the Boltzmann factors as our probabilities p_n and the system circuit will have a probabilistic mixture of the eigenstates $|\psi_n\rangle$ (once the ancillae are traced out). For distinct eigenvalues p_n , $|\psi_n\rangle$ are indeed energy eigenstates, but for degenerate eigenvalues, the density matrix eigenbasis can be a linear combination of energy eigenstates within that subspace. However, this is not a problem as the density matrix expressed in either basis (energy eigenbasis or a linear-combination basis) is still diagonal and represents the Gibbs states. We have added a comment in the main text explaining this within the paragraph that comes after Eq. (4) (on page 4, second column, beginning with 'At the end of the variational protocol...'). More mathematical details about how the thermal state is prepared can be found in the Methods Section E2 and is inspired from the more general construction given in Ref. [22] in the modified manuscript.

> **Reviewer 3:** It would be helpful if the authors provided some discussion on the importance of preparing thermal states in quantum simulators. In particular, it would be valuable to explain the challenges of incorporating finite-temperature effects in classical computations compared to zero-temperature calculations, and how quantum simulation may offer advantages in this regard. A comparison of the additional computational cost typically required, especially how this cost scales with system size beyond a single unit cell, in both classical and quantum approaches, would greatly enhance the context and significance of the present work. Additionally, it would strengthen the impact of the paper if the authors discussed whether their thermal state preparation method using motional ancillae is scalable to larger systems.

Our reply: We thank the reviewer for this insightful comment. While comparing classical and quantum resource requirements is crucial for demonstrating quantum advantage in thermal state preparation, our objective in this work is different. The experiment described in the paper is not aimed at showcasing an efficient method for preparing thermal states. Rather, it is an initial step toward addressing a longstanding open problem in physics: simulating the QCD phase diagram, using the latest advances in quantum computing.

Traditional classical methods (Lagrangian-based lattice QCD approach) of solving QCD fails in probing the phase diagram at finite chemical potential due to the sign problem. This renders a vast region of the $T - \mu$ phase diagram inaccessible to classical approaches. The quantum computer proves particularly advantageous in this aspect, as finite matter densities can easily be incorporated in this formulation, allowing us to probe the phase diagram at any chemical potential

values. This also requires us to create a thermal state in order to investigate the behavior of the order parameter in the entire $T - \mu$ plane, leading to a thermal state preparation protocol. So, our use of a quantum computer is motivated not by resource-efficiency over classical methods, but as a proof-of-concept for solving a problem that traditional classical methods cannot address.

It is of course worthwhile to mention that recent classical Hamiltonian-based methods involving tensor networks do not suffer from the sign problem. However, it is not known how to generalise the approach beyond one spatial dimension due to the exponential growth of the bond dimension with system size. For a quantum computer, generalising to higher spatial dimensions suffers no such fundamental issue, rather it depends on quantum hardware control capabilities and noise effects. As such, both quantum and tensor network methods represent complementary and exploratory avenues, with no clear consensus yet on which method is more advantageous. We have added a comment about this in the Discussion section (fourth paragraph, page 7) of the main text.

We would also like to emphasise the novelty and practical significance of the thermal state preparation method used in our experiment, focusing on its utility rather than its algorithmic efficiency. By using the motional modes as ancillae, we show that using ancilla-based protocols of creating thermal states does not necessarily mean increasing the system size. We can leverage these otherwise unused motional modes for quantum information processing in a thermal state preparation protocol. The use of the motional ancillae is thus not only relevant for our particular variational protocol, but will also be useful for future, more resource-efficient protocols relying on ancilla registers. Such approaches hold promise for simulating thermal states in broader quantum many-body contexts.

We also thank the reviewer for the question on scalability. The variational protocol we considered has two branches: the system circuit and the ancilla circuit. The circuit on the system is inspired from the terms present in the Hamiltonian. For a larger number of unit cells, we will have four-body and six-body nonlocal interactions in the SU(2) and SU(3) Hamiltonian, respectively, which can be implemented on a trapped-ion device due to the inherent all-to-all connectivity of the ions. The N-bodyness of the Hamiltonian does not grow with system size, which means the gates that appear in the circuit are also fixed. The number of MS gates needed to implement the Hamiltonian-inspired circuit thus only grows polynomially with the system size, making the circuit scalable. Having said that, for larger system size, we will have more parameters and optimisation becomes more challenging. So, we would eventually need a new optimiser to perform the global minimisation of the free energy.

Creating a probability distribution for a larger system will require entangling gates in the ancilla in general. This would require measuring the motional ancilla for entropy, which is challenging with the current hardware capabilities. Multiple methods have been used to measure phonon states in trapped ion systems. These methods are limited primarily by the motional coherence time of the phonon modes, such that only a small number of measurements can be performed on currently available devices. As the interest in using trapped ion systems for continuous variable quantum computing and bosonic quantum simulation is burgeoning, we expect these coherence times to improve significantly in the near future, enabling more measurements on larger devices. To clarify this in our manuscript, we now briefly describe the state-of-the-art and prospects for improving phonon measurement techniques in the section discussing the scalability of our method (Appendix E, section 2). However, the current single-qubit rotation architecture for ancilla circuit is still a good ansatz at certain regimes of chemical potential and temperature, making the ansatz scalable for larger system sizes.

We have added a paragraph in the Methods E2 (currently the second last paragraph in that section, page 11), which includes the above comments about scalability.

> **Reviewer 3:** Regarding the measurement method, it would be helpful if the authors discussed the general applicability of the approach presented in Eq. (5). For the specific case of the chosen order parameter—the chiral condensate in SU(2) and SU(3)—the expectation value can be evaluated using only qubit measurements in the computational basis. How-

> ever, it remains unclear how broadly this technique can be applied. For instance, can this projection-based measurement scheme be extended to systems with more lattice sites beyond a single unit cell? Is it applicable in higher-dimensional settings such as 2D or 3D lattice gauge theories? It would also be valuable to clarify to what extent this approach can serve as a substitute for explicitly implementing Gauss's law in the system.

Our reply: We thank the reviewer for this thoughtful comment. The projection-based approach is quite general and is applicable for projecting a state given in a reducible representation onto a particular irreducible representation. In our case, the density matrix on the full Hilbert space belongs to a reducible representation and we want to measure the observable on an irreducible singlet subspace. This description is thus not only limited to SU(2) and SU(3) LGT, but other gauge groups (including discrete gauge groups) as well.

The projection method in (1+1)-D is particularly simplified because the local gauge degrees of freedom can be integrated out to take care of Gauss's law. The job of the projection operator is then only enforcing the global color-neutrality or gauge-singlet condition, which is only dependent on the global color charges. For more than one unit cell, one just needs to insert the expression of the global color charge for multiple unit cell (given in Eq.(15) for SU(2) and Eqs. (21) and (26) for SU(3)) into the expression for the projection operator (Eq. (S5) for SU(2) and (S7) for SU(3)). Then since we are considering only diagonal charges, the integral needs to be performed for the diagonal entries of the integrand $e^{i\boldsymbol{\alpha}\cdot\hat{\boldsymbol{Q}}_{diag}}$. In fact, the closed form expression for the projection operator given in (S5) is valid for any number of unit cells.

For more than one spatial dimension, it is not possible to eliminate all gauge degrees of freedom. In that case, the projection operator can be generalised to enforce both Gauss's law on each vertex and the global color-neutrality condition. The form of the projection operator is more complicated and depends on the truncation scheme for the gauge fields. However, the statement that the effects of the gauge-constraints are implemented in the measurement stage still remains valid.

Other methods of implementing the global color-neutrality condition typically include adding penalty terms to the cost function or preparing a symmetry-preserving circuit. The penalty-based method becomes challenging for non-Abelian lattice gauge theories due to the presence of multiple non-commuting non-Abelian charges. This significantly complicates the optimisation landscape, making it difficult to reach the global minimum. On the other hand, constructing symmetry-preserving circuits is non-trivial for creating the gauge-singlet eigenstates. This challenge is further amplified in the context of thermal state preparation, which involves generating a probabilistic mixture of gauge-singlet states. This requires not only designing circuits that can prepare individual singlet states, but also ensuring that these states are combined with the correct Boltzmann weights. Simultaneously achieving both tasks is substantially more difficult than preparing a single gauge-singlet state. Our method provides an alternative in implementing the effects of the gauge-singlet condition as the task is delegated to the observable measurement stage. This allows us to employ any thermal state preparation protocol, involving more expressive circuits, without the need to explicitly impose color-neutrality during the state preparation itself.

We have added a sentence in the second paragraph of the Discussion section (page 7) that addresses the applicability of the method to higher dimensions.

In response to the comment 'For the specific case of the chosen order parameter—the chiral condensate in SU(2) and SU(3)—the expectation value can be evaluated using only qubit measurements in the computational basis', we want to mention that if we consider any non-diagonal observable O, the projection method and Eq. (5) will still be valid. However, we will need to measure non-diagonal Pauli strings that will appear in the decomposition of $\hat{O}\hat{K}$, in order to obtain $\langle\hat{O}\hat{K}\rangle$. So, only measuring the $z-$basis will not suffice anymore. We have added a comment on this in the main text in the paragraph that follows Eq. (5) on page 4.

Reviewer 3: When varying the chemical potential μ, what changes are required in the quantum simulation? Is μ simply treated as a parameter in the classical optimisation procedure for minimizing the free energy, or does it necessitate adjustments in the quantum circuit as well?

Our reply: We thank the reviewer for this question. The chemical potential enters the Hamiltonian as a free parameter, associated with the term $\hat{H}_{chem}$, whose explicit form is provided in Eq. (S12) for SU(2) and Eq. (S28) for SU(3) lattice gauge theories. Since $\hat{H}_{chem}$ is fully diagonal in the computational basis, μ appears only in the coefficients of certain diagonal Pauli strings in the Pauli decomposition of the Hamiltonian. As a result, varying μ affects only the post-measurement reconstruction of the Hamiltonian's expectation value–the circuit structure and measurement basis remain unchanged. However, the variational circuit parameters do depend on μ, since the contribution of each eigenstate to the thermal mixture is weighted by the Boltzmann factor $e^{-\beta(\tilde{E}_n - \mu B)}$, where B is the baryon number and $\tilde{E}_n$ is the energy at zero chemical potential. We have added a comment in the main text about the role of μ on the VQE ansatz at the end of the second paragraph of the quasi-subsection 'Experimental realisation' on page 5.

Reviewer 3: In Fig. 3, it appears that only a single experimental data point is presented at $\mu \approx 2.0$. Could the authors clarify why only this point was chosen? What are the main challenges in obtaining experimental results for other values of μ, and how do they differ from the SU(2) case where multiple data points were provided?

Our reply: We thank the reviewer for this question. The experimental implementation of the VQE for SU(3) is not fundamentally different from that of SU(2). In the SU(3) case, the variational optimization was carried out experimentally for only a single value of the chemical potential, whereas for SU(2), it was performed at five different points. This decision was driven purely by practical reasons—specifically, the limited experimental time available on the trapped-ion device.

The SU(3) VQE involves a greater number of variational parameters (21 compared to 10 for SU(2)) and requires control over more ions. As a result, both the calibration time and the number of function evaluations needed to reach a global minimum of the cost function increase. Given the limited run-time on the device, we chose to focus on a single chemical potential value for the experimental demonstration.

The value $\mu = 2.0$ was selected specifically to lie near the expected phase transition point. This choice was motivated by both physical relevance and experimental challenge, as accurately capturing behavior near a phase transition is typically the most demanding. Our rationale was that demonstrating successful VQE performance in this critical regime would provide strong evidence that the approach is expected to work at other values of the chemical potential as well. We have added a clarification in the third paragraph under the quasi-heading 'SU(3) thermal states' in the main text.

In addition, we would like to point out that the plot contains 15 experimental data points, 14 of which are based on direct implementation (shown in orange), which means the variational optimization is done classically before running the circuits on the trapped-ion quantum computer. For one point, the full VQE optimization was done on the quantum device as described above.

Reviewer 3: The variational circuit $\hat{U}_S$ shown in Fig. 1c seems relatively complex and resource-intensive. Beyond a brief mention of noisy numerical simulations, there is little discussion of the required gate fidelities and the actual performance of the quantum gates used in the experiment. How good were the quantum operations? What are the required fidelities for the method to remain accurate? Were any error mitigation techniques employed to

improve performance? A more detailed discussion of these aspects would be highly valuable.

Our reply: We agree that additional experimental details will improve the quality of our presentation. Alongside the spin-phonon gate description in Methods A (page 8), we have now added typical fidelities for the remaining quantum operations, with state detection fidelity discussed in the first paragraph and single- and two-qubit gate fidelities in the second paragraph. With these revisions, it is now clear in the manuscript that the entangling two-qubit Mølmer–Sørenson (MS) gates are responsible for the largest experimental error, motivating our use of the error model presented in the Methods section.

The minimum gate fidelities required for the experiment can be estimated using the error model outlined in Methods C (page 9). In the figure below, we illustrate how the target plot shown in Fig. 2c would be affected by varying the fidelity of the MS gates.

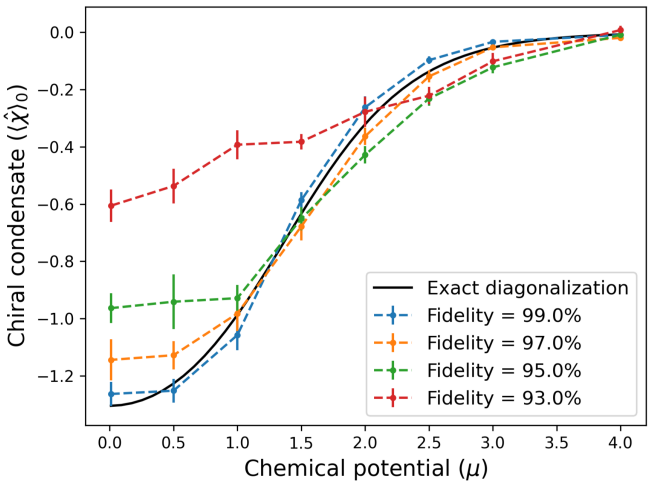

Figure 1: Results of noisy VQE optimization for varying fidelity of the MS gates. The quality of the captured transition of the chiral condensate value deteriorates as the noise level is increased.

The plot above shows that an MS gate fidelity exceeding 95% is sufficient for the device-aware noise model employed in our analysis. The experimentally achieved MS gate fidelity is above 98%, which lies above this threshold. However, it is important to note that the minimum required fidelity is not universal. It depends on the circuit size and depth, and will vary with the lattice size considered and the specific structure of the VQE ansatz. We have added a comment in the third paragraph of the Methods C (page 9) to accommodate this discussion.

By design, variational quantum algorithms mitigate experimental errors by optimizing the noisy experimental results. Hence, it is not necessary to implement costly mitigation techniques such as randomized compiling. We do, however, implement a simple method which largely eliminates state preparation and measurement error in the qubit basis. We have added details of this process to the Methods A (page 8), paragraph one.

We thank the reviewer again for providing valuable feedback to improve our manuscript and for his positive reviews of our work. We believe our revised manuscript addresses all the concerns above and will be of interest to the readers of Nature Communications.

Response to referee comments, 08/20/2025

We thank the reviewers for their time, careful evaluation, and insightful suggestions throughout the review process, which have significantly improved the clarity of our manuscript. We are pleased to learn that the article has been provisionally accepted. Here, we address the remaining comments raised by the reviewers and provide the corresponding clarifications in the revised manuscript.

. . . . . .

> **Reviewer 1:** The authors have satisfactorily addressed my previous comments and significantly improved the presentation, especially of the "physics application" component. With these changes, and given the innovation in the technical achievement of realizing motional ancillae, this is sufficient to justify publication in Nature Communications.
>
> I have one remaining recommendation. Although the authors have clarified the role of gauge symmetry in their model, they continue to use the term 'gauge-singlet', where 'charge-singlet' would be more appropriate. (Strictly speaking, in this work, the projection is applied to a global, not a local symmetry.) I don't mean to nitpick, but let me frame the concern constructively: consider the Heisenberg model, which has a global non-Abelian SU(2) symmetry. While this is not a gauge symmetry at all, the projection technique proposed by the authors might still be applicable. Referring to the relevant states as gauge-singlets rather than charge-singlets could unintentionally suggest a narrower scope of applicability than actually exists.

Our reply: We thank the reviewer for the positive assessment and for the helpful suggestion regarding terminology. We agree that charge-singlet is more appropriate than gauge-singlet in the context of our work. We have revised the manuscript accordingly, replacing all instances of gauge-singlet with charge-singlet throughout the text and figures.

. . . . . .

> **Reviewer 2:** I have reviewed the second version of this paper, together with the authors detailed answers to all three sets of referee comments. My original report was the most positive of the three and my original decision was to accept the paper subject to minor improvements. The authors have responded adequately to my questions and made small changes in the paper as a result.
>
> The other two referees are clearly more experienced in numerical methods and had many comments and criticisms of the method and models studied. Looking at the authors' answers and changes to the paper, I feel actually more confident than before in my decision to accept the paper. The main weakness of my report was indeed that I am no expert on numerical methods. As such my judgment was based on an appreciation of the importance of the problem that is being dealt with rather than a deep understanding of the methodology proposed. I am therefore very pleased that the other referees have been able to ask deeper questions about the methodology.
>
> My feeling is that the authors have dealt well with those questions and strengthen their paper as a result. Therefore I stand by my original decision to accept the paper for publication.

Our reply: We thank the reviewer for the appreciative remarks and for emphasizing the relevance and significance of the problem addressed in this work.

. . . . . .

> **Reviewer 3:** Overall, the authors have responded well, and I have no further objections to the publication of the paper. However, I would appreciate it if the following points could also be clarified:
>
> Regarding thermal state preparation, the current explanation still leaves some confusion. It gives the impression that one would need to know all eigenstates in advance, which would essentially mean the problem is already solved.
>
> In the fidelity figures, the case of SU(3) appears to show more serious deviations—would it be possible to provide additional data or discussion on this point? Also, since larger discrepancies occur in the low chemical potential region, I wonder if the experiment could have focused more on that regime.

Our reply: We thank the reviewer for the encouraging response and thoughtful questions. Below, we provide detailed clarifications in response to the points raised.

Q1: Regarding thermal state preparation, the current explanation still leaves some confusion. It gives the impression that one would need to know all eigenstates in advance, which would essentially mean the problem is already solved.

Our reply: We have revised the relevant section (page 4, first paragraph of the second column; now highlighted in blue) to clarify the explanation and address any potential confusion. In particular, we also added the following clarification: "Importantly, the basis states $|\psi_n\rangle$ are obtained through variational optimization of the cost function without requiring any prior knowledge of the energy eigenstates and the underlying spectrum".

The variational protocol described in the article prepares a thermal state by minimizing a free energy cost function. At the end of the VQE, the optimal circuit prepares a mixture over the states $\hat{U}(\boldsymbol{\varphi}^*)\,|j\rangle$, where the mixing probability is determined by the optimal parameters $\boldsymbol{\theta}^*$. Only if the VQE is successful in preparing the thermal state, $\hat{U}(\boldsymbol{\varphi}^*)\,|j\rangle$ would approximate the eigenbasis of the density matrix (not necessarily energy eigenbasis due to the degeneracy in the energy spectrum). Thus, the thermal state is prepared entirely through a variational procedure, without requiring any prior knowledge of the system's energy eigenstates.

Q2: In the fidelity figures, the case of SU(3) appears to show more serious deviations—would it be possible to provide additional data or discussion on this point? Also, since larger discrepancies occur in the low chemical potential region, I wonder if the experiment could have focused more on that regime.

Our reply: To address this question, we refer to Figs. 2c and 3c of the main text. For the SU(3) case, the number of variational parameters is significantly larger (21, compared to 10 for SU(2)), and the circuit involves a larger register and more complex structure. These factors lead to greater fluctuations during optimization, contributing to the observed variations in the chiral condensate. This is reflected in the larger grey boxes (representing noisy simulations) in Fig. 3c.

Notably, the grey boxes and error bars are largest in the intermediate region of the plot, near the phase transition. This aligns with our expectations that the transition region, where the system shifts from a vacuum-dominated to a baryon-dominated phase, is inherently harder to simulate reliably. To capture this physically relevant phenomenon and to test the performance of our protocol under challenging conditions, we focused our experimental efforts on this intermediate regime rather than the low chemical potential region.

We also emphasize that the red diamonds represent single experimental realizations of the full VQE procedure. As long as they lie within the error bars of the noisy simulation, this indicates consistency between experiment and theory, confirming that the VQE is functioning as intended.

We have added a comment at the end of page 6 (highlighted in blue) in the main text to incorporate our response.