nature portfolio

Peer Review File

Reviewers' comments:

Reviewer #1 (Remarks to the Author):

The manuscript by Inui and Motome presents a workflow for finding optimal Hamiltonians that maximize a predefined physical quantity. While the proposed flowchart that is standard for machine learning applications, the application to the inverse problem could be potentially new. My main concern is that the authors are presenting their workflow as an alternative to other methods, such as Bayesian optimization, genetic algorithms, and generative models and they do not show any comparison between the performance of their workflow and the performance of alternative methods. Without that comparison, it is not possible to judge if the proposed workflow is faster than the alternative methods.

The application of automatic differentiation to the optimization problem is very well-known by people working in artificial intelligence. Moreover, the quantum chemistry community has been using JAX for more than two years to obtain gradients automatically. See for instance:

https://neurips.cc/virtual/2020/public/poster_83d3d4b6c9579515e1679aca8cbc8033.html

Automatic Differentiation in Quantum Chemistry with Applications to Fully Variational Hartree-Fock - PubMed

DQC: A Python program package for differentiable quantum chemistry - PubMed

So, the only aspect of this manuscript that seems to be new is the application of JAX to the "Inverse Hamiltonian Design". While this may be enough to deserve publication in Communication Physics, I think that the authors should clarify their contribution in the introduction. If they want to claim that alternative methods have a higher computational cost, they should present evidence of that claim. The authors should also acknowledge that the flowchart depicted in Figure 1 is standard in Machine Learning applications and that their main contribution is to apply this flowchart to an optimization problem that is of interest for physicists. Another weakness of the presentation is that the authors do not talk much about the large literature on model inference.

Reviewer #2 (Remarks to the Author):

The authors propose a gradient-based approach for designing quantum materials by optimizing a desired physical property or observable. More specifically, they propose to calculate the gradient of an observable with respect to the Hamiltonian through automatic differentiation. This allows to scale the Hamiltonian design procedure to dozens of parameters in this manuscript, and potentially to much larger numbers in future works. They show this approach rediscovers the Haldane phase, and apply it to discover novel Hamiltonians with a particularly strong Hall conductivity and photovoltaic current, respectively.

To my understanding, the authors claim that previous approaches to quantum material design were

limited to a few free parameters, as they either (i) relied on model-specific analytical calculations, or (ii) used gradient-free ('black-box') methods which require prohibitively long times to converge when exploring large parameter spaces. The authors' automatic differentiation approach is model-agnostic, and allows to explore large parameter spaces much more efficiently compared to black-box approaches.
While I am not an expert in material science, Hamiltonian design by automatic differentiation seems to an outsider as me to be a natural, powerful and timely tool for engineering quantum materials, and I am not aware of previous works utilizing automatic differentiation for this particular purpose. The method proposed by the authors can be applied to a variety of models beyond those studied in this paper, and the particular experiments performed in the paper are well-motivated and thoroughly-analyzed. Their results are presented clearly and transparently, with publicly available code. Thus, I would recommend its publication in Communications Physics, subject to minor revisions.

Still, I think the paper can be improved in terms of its positioning within the existing literature. Firstly and most importantly, a more thorough discussion of the previous computational approaches for material exploration would be helpful. In particular, it would be helpful if the authors could elaborate which papers specifically were limited by each of the two limitations mentioned, i.e. high computational cost versus limited applicability. In addition, a more detailed analysis of the limitations and the computational complexity of the previous methods would allow to better appreciate the specific contribution of this paper.

Secondly, while not strictly necessary, I believe the paper would benefit from linking it to the broader community of quantum Hamiltonian engineering. For instance, in quantum computing, gradient-based methods using automatic differentiation are used used in optimal control schemes to design quantum gates – see e.g. [1,2]. Moreover, inverse Hamiltonian problems (sometimes called 'Hamiltonian learning') have recently gained interest in the theoretical condensed matter community; for instance, designing Hamiltonians with certain wavefunctions as their ground states [3,4].

Thirdly, an important (albeit perhaps obvious) limitation of the approach presented in the paper could be stated more clearly. Generally speaking, solving the inverse problem of Hamiltonian design through automatic differentiation has the same complexity as the forward problem of computing the desired property; therefore, it can only be extended to large system sizes for models which can be solved efficiently, such as the non-interacting systems studied in this paper. Note that there exist in the condensed matter community inverse Hamiltonian problems whose solution is much easier than the forward problem, but I am not sure they would be relevant to the problem studied here.

I attach a PDF with additional minor comments. In particular, some of the choices for parameterization of the Hamiltonian parameter space seem arbitrary at face value, and would benefit from further explanation, especially if these choices are important for future researchers wishing to apply the same methods to new models.

Refs:
[1] https://arxiv.org/abs/1612.04929
[2] https://arxiv.org/abs/1901.05541

[3] https://arxiv.org/pdf/1802.07827.pdf
[4] https://quantum-journal.org/papers/q-2020-09-02-315/

Summary of changes:

All the following changes are highlighted in blue in the revised manuscript.

1. In the first paragraph in page 1, we added sentences for the previous studies on the inverse approach and their limitations, citing the references including those suggested by Reviewers.
2. In the second paragraph in page 1, we added sentences for the previous studies using automatic differentiation in the field of physics, citing the references including those suggested by Reviewers.
3. In the third paragraph in page 1, we modified the sentence to state that the optimization flow is well-known in the machine learning.
4. In the fourth paragraph in page 1, we added sentences for the advantages of our framework in comparison with the previous studies. We also described the acronyms for AHE and PVE.
5. In Eq. (1), we added the summation by a, and changed a to a_i and d to d_{ij}. We also revised the text and the figure caption in pages 2 and 3 accordingly.
6. In the second paragraph in page 2, we revised the sentence to state the reasons for the parametrization.
7. In the first paragraph in page 4, we modified the sentences to state the motivation of this part.
8. In the first paragraph in page 4, we revised the sentence to explain the method to keep half filling. Also, we modified and added the sentences to clarify the optimization procedure.
9. In the third paragraph in page 5, we modified and added the sentences to explain the gist of this part.
10. In the third paragraph in page 5, we added sentences to state the reasons for the parametrization. We also added two references for the nonlinear optical conductivity for non-expert readers.
11. In the first paragraph in page 6, we revised the sentence to state the limitation of our framework and to make the description on the experimental raw data clearer.

In addition to these modifications, minor corrections were made for better readability.

Reply to Reviewer #1

We thank the reviewer for careful reading of our manuscript and pointing out the important issues. We answer the reviewer's comments below.

The manuscript by Inui and Motome presents a workflow for finding optimal Hamiltonians that maximize a predefined physical quantity. While the proposed flowchart that is standard for machine learning applications, the application to the inverse problem could be potentially new. My main concern is that the authors are presenting their workflow as an alternative to other methods, such as Bayesian optimization, genetic algorithms, and generative models and they do not show any comparison between the performance of their workflow and the performance of alternative methods. Without that comparison, it is not possible to judge if the proposed workflow is faster than the alternative methods.

We thank the reviewer for recognizing that our approach is potentially new and for providing valuable comments. We admit that our previous manuscript lacked comparison with other methods. We have revised the manuscript by providing comparisons with six algorithms: the perturbation theory, the potential interpolation, the eigenstate-to-Hamiltonian construction, the generative models, the Bayesian optimization, and the genetic algorithms (see the summary of changes 1). In addition, we added the sentences stating three advantages of our framework (see the summary of changes 4).

The application of automatic differentiation to the optimization problem is very well-known by people working in artificial intelligence. Moreover, the quantum chemistry community has been using JAX for more than two years to obtain gradients automatically. See for instance:
https://neurips.cc/virtual/2020/public/poster_83d3d4b6c9579515e1679aca8cbc8033.html
Automatic Differentiation in Quantum Chemistry with Applications to Fully Variational Hartree-Fock - PubMed
DQC: A Python program package for differentiable quantum chemistry - PubMed
So, the only aspect of this manuscript that seems to be new is the application of JAX to the "Inverse Hamiltonian Design". While this may be enough to deserve publication in Communication Physics, I think that the authors should clarify their contribution in the introduction. If they want to claim that alternative methods have a higher computational cost, they should present evidence of that claim. The authors should also acknowledge that the flowchart depicted in Figure 1 is standard in Machine Learning applications and that their main contribution is to apply this flowchart to an optimization

problem that is of interest for physicists. Another weakness of the presentation is that the authors do not talk much about the large literature on model inference.

We thank the reviewer for his/her constructive comments. We added a note on the previous studies using automatic differentiation including the reference that the reviewer pointed out (see the summary of changes 2). We also provided a comparison with the previous studies on computational costs and clearly stated the contributions of our framework (see the summary of changes 4). We revised a sentence to clearly state that the optimization flow in Fig. 1 is common in machine learning (see the summary of changes 3). With respect to model inference, we added descriptions of the previous studies (see the summary of changes 1).

Reply to Reviewer #2

We thank the reviewer for careful reading of our manuscript and recommending the publication. We answer the reviewer's comments below.

The authors propose a gradient-based approach for designing quantum materials by optimizing a desired physical property or observable. More specifically, they propose to calculate the gradient of an observable with respect to the Hamiltonian through automatic differentiation. This allows to scale the Hamiltonian design procedure to dozens of parameters in this manuscript, and potentially to much larger numbers in future works. They show this approach rediscovers the Haldane phase, and apply it to discover novel Hamiltonians with a particularly strong Hall conductivity and photovoltaic current, respectively.

To my understanding, the authors claim that previous approaches to quantum material design were limited to a few free parameters, as they either (i) relied on model-specific analytical calculations, or (ii) used gradient-free ('black-box') methods which require prohibitively long times to converge when exploring large parameter spaces. The authors' automatic differentiation approach is model-agnostic, and allows to explore large parameter spaces much more efficiently compared to black-box approaches.
While I am not an expert in material science, Hamiltonian design by automatic differentiation seems to an outsider as me to be a natural, powerful and timely tool for engineering quantum materials, and I am not aware of previous works utilizing automatic differentiation for this particular purpose. The method proposed by the authors can be applied to a variety of models beyond those studied in this paper, and the particular experiments performed in the paper are well-motivated and thoroughly-analyzed. Their results are presented clearly and transparently, with publicly available code. Thus, I would recommend its publication in Communications Physics, subject to minor revisions.

We appreciate the reviewer's comprehensive understanding and positive response to our work.

Still, I think the paper can be improved in terms of its positioning within the existing literature. Firstly and most importantly, a more thorough discussion of the previous computational approaches for material exploration would be helpful. In particular, it would be helpful if the authors could elaborate which papers specifically were limited by each of the two limitations mentioned, i.e. high computational cost versus limited applicability. In addition, a more detailed analysis of the limitations and the computational complexity of the previous methods would allow to better appreciate the

specific contribution of this paper.

We thank the reviewer for his/her constructive comment. We admit that our previous manuscript was insufficient in terms of the positioning within the existing literature. We have made revisions by adding comparisons with six methods: the perturbation theory, the potential interpolation, the eigenstate-to-Hamiltonian construction, the generative models, the Bayesian optimization, and the genetic algorithm (see summary of changes 1). We also added the sentences stating three advantages of our framework (see summary of changes 4).

Secondly, while not strictly necessary, I believe the paper would benefit from linking it to the broader community of quantum Hamiltonian engineering. For instance, in quantum computing, gradient-based methods using automatic differentiation are used used in optimal control schemes to design quantum gates – see e.g. [1,2]. Moreover, inverse Hamiltonian problems (sometimes called 'Hamiltonian learning') have recently gained interest in the theoretical condensed matter community; for instance, designing Hamiltonians with certain wavefunctions as their ground states [3,4].
[1] https://arxiv.org/abs/1612.04929
[2] https://arxiv.org/abs/1901.05541
[3] https://arxiv.org/pdf/1802.07827.pdf
[4] https://quantum-journal.org/papers/q-2020-09-02-315/

We thank the reviewer for his/her useful remarks. We have added a note on quantum Hamiltonian engineering and Hamiltonian learning, citing the four references (see the summary of changes 1 and 2).

Thirdly, an important (albeit perhaps obvious) limitation of the approach presented in the paper could be stated more clearly. Generally speaking, solving the inverse problem of Hamiltonian design through automatic differentiation has the same complexity as the forward problem of computing the desired property; therefore, it can only be extended to large system sizes for models which can be solved efficiently, such as the non-interacting systems studied in this paper. Note that there exist in the condensed matter community inverse Hamiltonian problems whose solution is much easier than the forward problem, but I am not sure they would be relevant to the problem studied here.

The reviewer is right: there are limitations when the forward problem is computationally expensive, for example, in quantum many-body problems. We added a comment on this point (see the summary of changes 11).

Regarding the last point, we are not familiar with the faster method for the inverse problem. We would appreciate it if the reviewer could provide the references.

I attach a PDF with additional minor comments. In particular, some of the choices for parameterization of the Hamiltonian parameter space seem arbitrary at face value, and would benefit from further explanation, especially if these choices are important for future researchers wishing to apply the same methods to new models.

We appreciate the reviewer's careful reading and detailed comments. Carefully examining the comments, we revised the manuscript (see the summary of changes 1, 4-11).

REVIEWERS' COMMENTS:

Reviewer #1 (Remarks to the Author):

The authors have improved the presentation of their results based on the input from both reports. The new presentation provides a more general context that allows the reader to understand the value of the contribution. I recommend the new version of the manuscript for publication in Communications Physics.

Reviewer #2 (Remarks to the Author):

All the points raised in the previous round of reviews have been addressed.
The newly-added comparison to other methods could be improved in terms of clarity and quality of writing. Still, the current comparison is more comprehensive compared to the original version.

(Regarding computationally-efficient inverse methods for many-body problems mentioned in the previous review: I was referring to methods such as the eigenstate-to-Hamiltonian approach which is now cited, as well as other methods which have been recently proposed in the condensed matter and quantum computation communities for learning Hamiltonians from their steady states and thermal states; however, I am not sure they are directly applicable to the problem studied here)

Overall, the paper is now improved in terms of its positioning in the literature as well as clarity due to many changes throughout the manuscript

Reply to Reviewer #1

We thank the reviewer for careful reading of our manuscript and supporting the publication.

The authors have improved the presentation of their results based on the input from both reports. The new presentation provides a more general context that allows the reader to understand the value of the contribution. I recommend the new version of the manuscript for publication in Communications Physics.

We thank you for acknowledging our revised paper. We also thank you for recommending our paper for publication.

Reply to Reviewer #2

We thank the reviewer for careful reading of our manuscript and recognizing the improvement of our revised manuscript.

All the points raised in the previous round of reviews have been addressed.
The newly-added comparison to other methods could be improved in terms of clarity and quality of writing. Still, the current comparison is more comprehensive compared to the original version.

(Regarding computationally-efficient inverse methods for many-body problems mentioned in the previous review: I was referring to methods such as the eigenstate-to-Hamiltonian approach which is now cited, as well as other methods which have been recently proposed in the condensed matter and quantum computation communities for learning Hamiltonians from their steady states and thermal states; however, I am not sure they are directly applicable to the problem studied here)

Overall, the paper is now improved in terms of its positioning in the literature as well as clarity due to many changes throughout the manuscript

We thank the reviewer for recognizing that the comparison to other methods is more comprehensive than the previous manuscript. We also appreciate the acknowledgment of the clarity of our paper.