Response to Reviewer 1

Summary Comment

The authors have addressed a number of points but significant concerns remain. As elaborated below, chief among these concerns are the opacity of the presentation of the proposed RCG and RBFGS methods, despite multiple reviewer comments to clarify it, and the lack of fairness in comparisons to the standard SCF method, which may therefore be misleading. Only if the significant concerns below can be fully addressed, would I recommend publication.

Response:

We would like to express our gratitude again for your valuable suggestions on our manuscript. We have carefully reconsidered each of your points and reflected upon ourselves. Below are point by point responses to your concerns.

Comment 1.1

Abstract: "Extensive testing compares the performance of the proposed methods with the traditional self-consistent field (SCF) algorithm and shows that it is less efficient." This statement relies on comparisons to the standard SCF method run in nonstandard way, and so may be significantly misleading.

As noted in the previous review, at a minimum, a Kerker-type preconditioner is typically employed in SCF methods in cases with significant partial orbital occupations, as occur for metallic systems, and it is *not* typically employed for isolated systems (though it may be tried in cases with partial occupations and slow SCF convergence). Furthermore, when SCF convergence is slow or fails, the mixing parameter ("mixing_beta" in the ABACUS software the authors use) is typically decreased. Indeed, the ABACUS manual https://abacus.deepmodeling.com/en/latest/advanced/scf/converge.html states: "For most isolated systems, Kerker preconditioning is unnecessary. You can turn off it by setting mixing_gg0 0.0 to get a faster convergence." and "For non-spin-polarized calculations, the default choices usually achieve convergence. If convergence issue arises in metallic systems, you can try different value of Kerker preconditioning mixing_gg0 and mixing_gg0_min, and try to reduce mixing_beta, which is 0.8 defaultly for nspin=1." But rather than follow standard practice, the authors instead just use all defaults for the vast majority of the systems considered. Thus, the comparisons to the SCF method are not representative of the SCF method in practice, and so may be significantly misleading. This must be fully addressed. Preferably by using/not-using Kerker preconditioning and adjusting the mixing parameter in the SCF calculations consistent with standard practice. Alternatively, if the authors deem it too much effort to provide such fair comparisons to the standard SCF method, they should at the very least clearly explain the limitations of the comparisons they present, due to not following standard practice in setting the most basic parameters. And if opting to explain rather than address the unrepresentative nature of the comparisons, all assertions of relative efficiency – in both iterations and time to solution – should be removed, since they are with respect to unrepresentative SCF calculations. (Another issue is the use of different stopping criteria for SCF and RCG/RBFGS methods. More on that below.)

Appendix B: In the figure, the convergence criterion is in terms of a difference of densities rather than a *relative* difference of densities, as discussed on p. 10. Again, a single convergence criterion should be chosen and used for all methods being compared. Otherwise, comparisons are not meaningful.

Response 1.1:

Thank you for insisting on the fair comparison. Given the differences in

- main algorithm

- choice of conditioner
- choice of mixing parameters
- tolerances for convergence

it is indeed too much effort to provide a stringent comparison to the SCF method. To avoid possible misunderstandings, we have decided to remove the comparison on iterations and time against SCF. It is not our intention to challenge the widely-adopted iterative diagonalization practice, while we believe the efficiency and robustness has to come later through continous effort on refinement of the implementation.

The convergence criterion in the flowchart is in fact one choice. It could be either the relative density error or the energy difference or other variants. Our comparison with SCF is performed to ensure the correctness of the proposed methods, namely the converged energy agrees with that from SCF for a given tight error tolerance. We hope this clarifies the situation.

- We have removed on iterations and time related content on SCF.
- We have added a note on further effort, which now reads "We anticipate that the efficiency and robustness of direct minimization methods can be further enhanced through continous effort on refinement of the implementation. "

Comment 1.2

p. 2: Clarification of what is actually new/significant in the present work: "We note that a recent 2023 paper [34], which is motivated by metallic systems and employs a plane-wave basis [34], aligns closely with our approach in this regard, despite our use of a localized basis in the current work." In what "regard"? "despite" our use? This statement needs clarification. For example, if I've understood the authors' response correctly, something like the following would be much clearer: "We note that a recent 2023 paper [34], which is motivated by metallic systems and employs a plane-wave basis [34], uses Bloch-periodic boundary conditions for periodic calculations, as we do here; whereas our formulation is in terms of a general, non-orthogonal basis rather than a global, orthonormal one as in Ref. [34].

Response 1.2: Thank you for your suggestion on the presentation.

We have made such corrections.

Comment 1.3

p. 4: "(bold symbol for this size)" What is meant by "size"? What "size" corresponds to a bold symbol? What "size" corresponds to a non-bold symbol? Please clarify.

Response 1.3:

The size of $M \times N \times K$ is meant here, as explained in Section 2.1. A non-bold symbol has size of $M \times N$ or the special case of $K = 1$. We have made the clarification in the revised version. Now it reads "(bold symbol for this size, see notations in Section 2.1)".

Comment 1.4

p. 7, Eqs. (41a) and following: $<,>_{X_k}$ notation is not defined.

Response 1.4: Thank you for pointing out this omission.

We have refined the definition of canonical metric in Eq. 32 (now Eq. 33) with a subscript of X to make it clear.

Comment 1.5

p. 7: The authors now partially address the opacity of their functional notation by providing essentially a quote of my best guess as to how it translates to the actual matrix operations performed in one case (the "retraction" operation). And then follow with, "For other retractions and vector transports, one needs to apply the corresponding linear algebra operations." That's correct. The problem is (1) I had to *guess* the corresponding linear algebra operations for the retraction function, and (2) I would have to guess the others as well. As commented in the previous review, the authors should provide the corresponding linear algebra operations for *all* retractions and vector transports occurring in all algorithms and figures (Algorithm 1, Fig. 1, and Algorithm 2 in the present version). Otherwise, it is simply not clear what the proposed algorithms are actually doing, and so not possible for the reader to understand, implement, compare to previous works, and/or assess.

Response 1.5: Thank you!

To further elaborate these operations, we have made a new table for related linear operations and pertinent key LAPACK routines in the revised version.

Comment 1.6

p. 7: Regarding k-points, the authors now say, "To apply to the multiple k-points cases, one simply expands the dimension of the pertinent X, by stacking K copies of X, each of the same size (see below the product manifold)." This is helpful but still not enough. E.g., is the minimization over the entire set of K copies (unlikely) or over each X_i separately? To ensure all is clear, k-points should be included in Algorithm 1, Fig. 1, and Algorithm 2.

Response 1.6:

It means that the minimization should be done over the entire set of K copies (bold symbol). We have made corresponding changes to clarify this point. Now it reads:

> For other retractions and vector transports, one needs to apply the corresponding linear algebra operations. To apply to the multiple **k**-points cases, one simply expands the dimension of the pertinent X, by stacking K copies of X, each of the same size. In the algorithm, one needs to adapt X and g into bold symbols $\mathbf{X}$ and $\mathbf{g}$ (see below the product manifold).

Comment 1.7

p. 8: The stopping criterion in Algorithm 1 is different than that in the corresponding Figure 1. Whatever stopping criterion was actually used to produce the results presented should be used consistently in all algorithms and associated figures.

Response 1.7: Thank you for your suggestion.

We have modified the stopping criterion in Algorithm 1 to allow for more choices.

Comment 1.8

Fig 1: Should $H_s[n]$ be $H_s[n_{k+1}]$?

Response 1.8: Thank you for the comment.

$H_s[n]$ is meant to take the current density n as input to form the Hamiltonian, which is $H_s[n_{k+1}]$. We have made the corrections in the revised version.

Comment 1.9

p. 10: "The nonlocality of pseudopotential is included via the standard nonlocal projectors [39]." How is the nonlocal term taken into account in the derivatives of the energy functional (Eqs. 27 & 28)? For clarity, the nonlocal part should be included in the mathematical formulation, from start to finish. The reader should not have to guess how to incorporate such significant terms into the mathematical formulation.

Response 1.9: Thank you for your valuable comments.

We have explicitly explained the nonlocal terms, which is only relevant in the construction of Hamiltonian. There is no need to change Eqs 27 and 28.

Comment 1.10

p. 10: Regarding the SCF method, it is stated that the convergence criterion is that the relative density error is less than 10^{-6}. While for RCG and RBFGS the convergence criterion is that 2 successive energies differ by less than 10^{-8} Ry. This is highly problematic for two reasons. First, since the convergence criteria are different for SCF and RCG/RBFGS methods, comparisons of the number of iterations and/or time to solution for SCF and proposed RCG/RBFGS methods are not meaningful. In order to have fair and meaningful comparisons of iterations or time to convergence, the convergence criteria must be the same for the methods compared. Since either criterion is simple to insert in any codes being compared, there is no good reason not to. Second, the criteria do not scale the same with system size. That is, while relative density error is independent of system size, difference of total energies is proportional to system size. So the different criteria are fundamentally incompatible for comparison purposes. (Also, why is the energy criterion in units of Rydberg rather than Hartree as in the rest of the manuscript?)

Response 1.10:

We agree that the absolute energy difference is less optimal compared to intensive measures when system sizes vary. The comparison to the SCF method now is to ensure the correctness of the RCG and RBFGS methods for a given error tolerance. That is, for each system, the converged energy from Riemannian methods should be within the error tolerance of the SCF method. The unit of Rydberg used here literally follows the code implementation, as we have used ModuleBase::Ry_to_eV in the print_info function. We have now changed it to Hartree in the revised version.

Comment 1.11

p. 11: "Calculations for all 148 molecules in the G2 dataset are successfully converged using the SCF method. However, for some molecules, the minimization process of the RCG method halted after the first few iterations," Since a main motivation for the minimization methods was robustness relative to the standard SCF method, it would be good, if at all possible, to figure out why the RCG method only converges for 134 of the 148 molecules considered while the SCF method converges for all 148. And remedy the issue in RCG so that it too converges for all 148. Otherwise, the results demonstrate the opposite of what is stated with respect to motivation and robustness.

Response 1.11:

We are currently adapting the code to the develop branch, the early termination of the RCG

method is still under investigation. We have been working on adding more unit tests to each of the components to ensure the robustness of the implementation.

Comment 1.12

p. 12: "the diagonalization step in SCF or the manifold-related operations in RCG account for only a small fraction" Which operations exactly are the "manifold-related operations"? Please clarify.

Response 1.12: We sincerely appreciate your valuable suggestions.

They are the operations related to the manifold, such as the retraction and vector transport operations, as well as metric evaluations.

Comment 1.13

Appendix A: What is new here relative to previous works [26,60,61,62]?

Response 1.13:

Essentially, this appendix mostly follows the work of [61,62], and the other references are general.