nature research

Peer Review Information

Journal: Nature Computational Science
Manuscript Title: Efficient sampling of high-dimensional free energy landscapes using adaptive reinforced dynamics
Corresponding author name(s): Linfeng Zhang

Reviewer Comments & Decisions:

Decision Letter, initial version:

Dear Dr Zhang,

Thank you for submitting "Efficient sampling of high-dimensional free energy landscapes using adaptive reinforced dynamics" to Nature Computational Science. Regretfully, we cannot offer to publish it in its current form.

Among the considerations that arise at this stage are the manuscript's likely interest to a broad range of researchers in computational science, the pressure on space for the various disciplines covered by Nature Computational Science, and the likelihood that a manuscript would seem of great topical interest to those working in the same or related areas of computational science. We do not doubt the technical quality of your work or that it will be of interest to others working in this area of research. However, I regret that we are unable to conclude that the paper provides the sort of substantial practical or conceptual advance that would be of immediate interest to a broad readership of researchers in computational science.

Should future experimental data allow you to address the following points, we would be happy to look at a revised manuscript (unless, of course, something similar has by then been accepted at Nature Computational Science or appeared elsewhere). This includes submission or publication of a portion of this work somewhere else. In the case of eventual publication, the received date would be that of the revised paper.

- Please better clarify the methodological novelty of the presented work over your earlier RiD work [J. Chem. Phys. 148, 124113 (2018)], such as what has been made possible by this development?

- Please provide data comparisons against other methods, such as metadynamics and REMD, to demonstrate your claimed benefits of your modified version of RiD .

If you are interested in submitting a suitably revised manuscript in the future or if you have any questions, please contact me.

Thank you for your interest in Nature Computational Science. I am sorry that on this occasion we cannot be more positive.

Best regards,

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

Author Rebuttal to Initial comments

Dear Dr. Pan,

Enclosed please find the manuscript entitled “Efficient sampling of high-dimensional free energy landscapes using adaptive reinforced dynamics” that we are pleased to submit to Nature Computational Science. After a brief communication with you, we believe that the revised manuscript has addressed well the issues you raised, and we are confident that the current manuscript will meet the publication standard of Nature Computational Science. We have prepared a version of our manuscript that highlights all the changes and sentences that are related to the questions you raised in our earlier discussion, and we have also uploaded a version of the main manuscripts without highlights and our Supplementary Materials.

Enhanced sampling methods such as metadynamics and replica exchange schemes have become essential tools for exploring the configuration space of molecules and materials. However, for several challenging problems, such as ab initio protein folding, protein structure refinement, protein-peptide interaction, etc., there have been essential difficulties for these schemes to be made routine procedures. In other words, human insights and human interventions still play important roles in these schemes, and their efficiencies are not satisfactory.

In the methodology that we propose here, we employ a large number of collective variables (CVs) to help efficiently explore the configuration space. Deep neural networks (DNNs) are used to represent the free energy of these CVs and an adaptive procedure is used to optimize the DNNs and enhance the sampling process. This is a procedure that is almost fully automatic, and we illustrate its usefulness using

three challenging examples. In particular, as far as we can see, there has not been an effective enhanced sampling procedure for the example of protein refinement.

As in the previous discussion with you, we notice that the current scheme is an extension of our earlier RiD work [J. Chem. Phys. 148, 124113 (2018)]. Compared with our earlier work, various new features of the current methodology, such as the adaptive error thresholds, the clustering techniques, as well as the multi-walker support, make possible an enhanced sampling technique truly driven by more than 100 CVs. Our numerical examples can well support this. In particular, we provided a detailed analysis of this using the first example (peptoid). We didn't use the original version of RiD for the second and the third examples (protein folding and protein structure refinement, respectively), since from the first example we can conclude that the current methodology will have a much better performance.

In addition, we have tried our best to compare our methodology with existing ones using our examples. We prepared detailed analysis and data comparisons for the first example. For the second and the third example, we didn't perform additional simulations using either metadynamics or REMD, but referred to existing works and made comparisons. We found strong support that methodologies like metadynamics can hardly work in these cases, since the number of collective variables involved is too large; on the other hand, REMD was sometimes used, but it is much less efficient than our methodology and has some additional drawbacks. We also compared with some recent schemes like variationally enhanced sampling (VES), which requires explicit prior knowledge of the system and is much less efficient. We've made all these considerations more explicit in our revised manuscript.

As possible referees for this work, we would like to suggest Jianfeng Lu (Duke University, jianfeng@math.duke.edu), Weiliang Zhu (Shanghai Institute of Materia Medica, wlzhu@simm.ac.cn), Eric Vanden Eijden (New York University, eve2@cims.nyu.edu), and Yiqin Gao (Peking University, gaoyq@pku.edu.cn).

We thank you in advance for your consideration. Please feel free to contact me, via linfeng.zhang.zlf@gmail.com, as the corresponding author to discuss any questions you may have about this work or this manuscript.

Kind regards,

Linfeng Zhang

Program in Applied and Computational Mathematics

Princeton University, Princeton, NJ 08544

Decision Letter, first revision:

Dear Dr Zhang,

Thank you again for submitting your manuscript "Efficient sampling of high-dimensional free energy landscapes using adaptive reinforced dynamics". I am pleased to tell you that we are sending your paper out for formal peer review. Before we can do so, please read the below carefully as we require a few documents.

If you have not done so already, please alert us to any related manuscripts from your group that are under consideration or in press at other journals, or are being written up for submission to other journals (see https://www.nature.com/authors/policies/duplicate.html for details).

We are asking all corresponding authors of primary research articles to complete an Editorial Policy Checklist that verifies compliance with all required editorial policies. Please note that the form is a dynamic 'smart pdf' and must therefore be downloaded and completed in Adobe Reader. We will then flatten them for ease of use by the reviewers. If you would like to reference the guidance text as you complete the template, please access these flattened versions at https://www.nature.com/authors/policies/availability.html

Editorial Policy Checklist: https://www.nature.com/documents/nr-editorial-policy-checklist.zip

In addition, as your paper relies on code that is central to the main claims, we will ask the reviewers to evaluate the code during the peer review process (for more details on this please see this editorial in Nature https://www.nature.com/articles/d41586-018-02741-4). In this case, it would be the code that implements your adaptive reinforced dynamics framework, and the code/data that reproduces your results/figures in the manuscript.

Reproducibility and re-usability of code are very important to us so, to facilitate this process, we are currently running a trial in partnership with Code Ocean to enable authors to share fully-functional and executable code accompanying their articles and to facilitate peer review of code by the reviewers (for more details please see http://blogs.nature.com/ofschemesandmemes/2018/08/01/nature-research-journals-trial-new-tools-to-enhance-code-peer-review-and-publication). We expect this functionality to speed up the peer review of your paper as it will facilitate the reviewer's assessment.

The use of the Code Ocean platform for peer review of code associated with this paper will be under the same confidentiality and anonymity agreements as the rest of manuscript materials.

Code Ocean is a cloud-based reproducibility platform where authors upload code and data and configure the necessary computational environment for reproduction. The code, data, metadata, and computational environment -- called a 'compute capsule' -- can then be accessed by reviewers in an anonymous fashion, and upon publication, provided to readers via a link from the article. Code Ocean supports all open source programming languages, as well as Stata and MATLAB and compute capsules can be created from existing GitHub folders by easy drag and drop.

Code Ocean staff will assist you in generating a compute capsule for your code and a working copy of this compute capsule will be used in the peer review process (after a brief review by Code Ocean staff, to ensure that everything runs). If you have selected Double Blind Peer Review, we will make sure the capsule contains no information about your identity.

If your code is accepted for publication, Code Ocean will assign a Digital Object Identifier (DOI) to your compute capsule. It will then be embedded into your article. Code Ocean, through CLOCKSS, will guarantee the preservation of all elements of Code Ocean's compute capsules, including the code, data, results, metadata, Dockerfile and Docker image (computational environment) associated with your paper.

By using this platform, other researchers will then be able to easily find and run the code, as well as build upon your work, without any additional setup or configuration of the software. It will also enable preservation of your code, data and the complete environment so that the code associated with this publication is maintained. Should the paper be rejected, you will retain full control over the compute capsule, and be able to decide what to do with it (publish it, modify it etc).

We very much hope you will be interested in engaging in this trial. Please let us know as soon as possible if you wish to participate and we will provide you with further guidelines for setting up the compute capsule. An overview of the process can be found here: https://help.codeocean.com/publishing-on-code-ocean/peer-review-on-code-ocean .

Alternatively, if you do not want to engage in this trial, or if Code Ocean is not a good fit (https://help.codeocean.com/en/articles/3294415-what-is-and-is-not-a-good-fit-for-publishing-and-sharing-on-code-ocean), we ask you to complete the following Software and custom code submission checklist:

Software supplement: https://www.nature.com/documents/nr-software-policy.pdf
(Please note that the form is a dynamic 'smart pdf' and must therefore be downloaded and completed in Adobe Reader.)

To improve transparency in authorship we are requesting that all authors identified as 'corresponding

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

Finally, we encourage you to share a preprint of the original submitted version of your paper so as to minimize delays in communicating your research findings; benefits of preprints include early visibility, and citations (https://www.natureindex.com/news-blog/preprints-boost-article-citations-and-mentions) and demonstration of research progress. You may want to consider the multidisciplinary Research Square preprint platform (https://www.researchsquare.com/browse), provided by our partner Research Square, where your preprint will be publicly available with a citable DOI under a CC-BY license. You are of course free to use a discipline-specific preprint platform of your choice. More information about our preprint policy can be found in the following link: https://www.nature.com/nature-research/editorial-policies/preprints-and-conference-proceedings#preprints

Please use the following link to submit the required checklists; please also resubmit your original manuscript files, or revised versions of them as a result of filling out the checklists:

[***REDACTED***]

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

Thank you very much for your attention to this. We look forward to hearing from you, by replying to this email, about the Code Ocean trial and any related manuscripts, but please let me know if you have any questions.

Best regards,

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

Author Rebuttal, first revision:

Dear Dr. Pan,

Enclosed please find the manuscript entitled “Efficient sampling of high-dimensional free energy landscapes using adaptive reinforced dynamics” that we are pleased to submit to Nature Computational Science. After a brief communication with you, we believe that the revised manuscript has addressed well the issues you raised, and we are confident that the current manuscript will meet the publication standard of Nature Computational Science. We have prepared a version of our manuscript that highlights all the changes and sentences that are related to the questions you raised in our earlier discussion, and we have also uploaded a version of the main manuscripts without highlights and our Supplementary Materials.

Enhanced sampling methods such as metadynamics and replica exchange schemes have become essential tools for exploring the configuration space of molecules and materials. However, for several challenging problems, such as ab initio protein folding, protein structure refinement, protein-peptide interaction, etc., there have been essential difficulties for these schemes to be made routine procedures. In other words, human insights and human interventions still play important roles in these schemes, and their efficiencies are not satisfactory.

In the methodology that we propose here, we employ a large number of collective variables (CVs) to help efficiently explore the configuration space. Deep neural networks (DNNs) are used to represent the free energy of these CVs and an adaptive procedure is used to optimize the DNNs and enhance the sampling process. This is a procedure that is almost fully automatic, and we illustrate its usefulness using three challenging examples. In particular, as far as we can see, there has not been an effective enhanced sampling procedure for the example of protein refinement.

As in the previous discussion with you, we notice that the current scheme is an extension of our earlier RiD work [J. Chem. Phys. 148, 124113 (2018)]. Compared with our earlier work, various new features of the current methodology, such as the adaptive error thresholds, the clustering techniques, as well as the multi-walker support, make possible an enhanced sampling technique truly driven by more than 100 CVs. Our numerical examples can well support this. In particular, we provided a detailed analysis of this using the first example (peptoid). We didn't use the original version of RiD for the second and the third examples (protein folding and protein structure refinement, respectively), since from the first example we can conclude that the current methodology will have a much better performance.

In addition, we have tried our best to compare our methodology with existing ones using our examples. We prepared detailed analysis and data comparisons for the first example. For the second and the third example, we didn't perform additional simulations using either metadynamics or REMD, but referred to existing works and made comparisons. We found strong support that methodologies like metadynamics can hardly work in these cases, since the number of collective variables involved is too large; on the other hand, REMD was sometimes used, but it is much less efficient than our methodology and has some additional drawbacks. We also compared with some recent schemes like variationally enhanced sampling (VES), which requires explicit prior knowledge of the system and is much less efficient. We've made all these considerations more explicit in our revised manuscript.

As possible referees for this work, we would like to suggest Jianfeng Lu (Duke University, jianfeng@math.duke.edu), Weiliang Zhu (Shanghai Institute of Materia Medica, wlzhu@simm.ac.cn), Eric Vanden Eijden (New York University, eve2@cims.nyu.edu), and Yiqin Gao (Peking University, gaoyq@pku.edu.cn).

We thank you in advance for your consideration. Please feel free to contact me, via linfeng.zhang.zlf@gmail.com, as the corresponding author to discuss any questions you may have about this work or this manuscript.
Kind regards,

Linfeng Zhang
Program in Applied and Computational Mathematics
Princeton University, Princeton, NJ 08544

Decision Letter, second revision:

Dear Dr Zhang,

Your manuscript "Efficient sampling of high-dimensional free energy landscapes using adaptive reinforced dynamics" has now been seen by 3 referees, whose comments are appended below. You will see that while they find your work of interest, they have raised points that need to be addressed before we can make a decision on publication.

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

While we ask you to address all of the points raised, the following points need to be substantially worked on:

- As suggested by referee #2 and #3, please provide an improved discussion about novelty/technical improvements when compared with the most recent developments in the fields, e.g., other machine learning approaches, metadynamics methods, and your recent work on the topic [The Journal of chemical physics, 148(12):124113, 2018]
- Referee #2 suggested that a better benchmark/comparison study in efficiency is needed. Specifically, this referee suggested to benchmark the efficiency against Bias-Exchange MetaD and/or PBMetaD using the peptoid and chignolin test cases with 3 and 18 CVs, respectively.
- Referee #1 suggested to be more detailed about the conversion of a multi-dimensional potential

energy surface to low dimensional PES (this will be important to predict variables to be compared with experimental observations)
- Referees have concerns about the current presentation of technical points. Please clarify those points as suggested, such as a better discussion about your deep neural network (referee #1 and #3 have concern about the deepness of your framework), the details about whether TAMD is used, etc.
- Jargons, like GDT-HA score as used in the protein structure prediction and refinement field, need to be better defined for our general audience. Please define jargons or specific scientific terms appropriately in your manuscript.
- Please better discuss the generality and the limitations of your method. For example, referee #3 has concerns about point that how general your approach is for the treatment of different type of CVs.
- Please provide better justification on your neural network structure; for instance, was an ablation study performed? How researchers from other fields can learn from your neural network model? These detailed information will be important to our multidisciplinary readership and also increase the potential impact of your paper.

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

[***REDACTED***]

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

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

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

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

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

We look forward to hearing from you soon.

Best regards,

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

Reviewers comments:

Reviewer #1 (Remarks to the Author):

The manuscript of Wang and coworkers presents a new variant of the method for sampling enhancement in molecular simulations. It uses a neural network (I cannot judge whether neural networks used here are enough deep to be called “deep”) to approximate a high-dimensional free energy surface of the simulated system. This estimate is used to bias the simulation. Biasing is adaptively switched on or off based on uncertainty of the free energy surface estimate. The method was tested on three types of molecular systems differing in complexity and heights of energy barriers. Some of the tests systems are challenging. This makes the manuscript suitable for publication after some improvements:
1. It is not clear how low-dimensional free energy surface can be obtained from high dimensional one. As far as I understand the method, the output is a high-dimensional FES. For practical applications, such as prediction of protein thermostability or binding free energy of a protein-ligand complex, it is useful to have a 1D or 2D FES, not e.g. 20D. This is because we want to predict some variable that can be compared with an experiment. Conversion of 20D FES to 1D FES may be a nontrivial task. At the beginning I suspected this to be a weak point of the method and I suspected that this was the reason why authors focused on protein structure refinement, where this problem is less important. Later I found Fig. 2 where 1D FESes are presented, so the method can predict lowdimnesional FESes. I would ask authors to give more details on conversion of a multidimensional FES to a lowdimnesional.
2. The article is likely to be read mostly by molecular simulation experts. The authors use GDT-HA score without any explanation, even without explaining the abbreviation. Authors should add one or two sentences introducing GDT-HA to nonexpert in protein structure predictions. This should be addressed also in the abstract. The abstract should be self explanatory.
3. The switching function sigma is not clear to me. I understood that it is sigma(epsilon) = 0 for epsilon < e0, sigma(epsilon) = 1 for epsilon > e1, and sigma(epsilon) = (epsilon-e0) / (e1-e0) elsewhere. Is this

correct? I would appreciate more clear explanation.
4. Evolution of e0 and e1 is quite confusing in Figure 1 and the related text. I would advice authors to use some symbol for iteration number, e.g. “t” (of course authors may use other symbol or notation) to replace e0=e0 by e0(t) = e0(0) and to replace e0 = e0 * 1.5 by e0(t) = e0(t-1) * 1.5 (maybe I’m wrong with the use of t, t-1 and 0, but this highlights the need for more clear mathematical description).
5. I appreciate that authors provide the code at GitHub. As far as I can judge it is well established. I advise authors to also submit their input files via Plumed Nest. I understand that there might be some obstacles because Plumed has been modified for this work. Regarding the modification, would it be possible to add it to official Plumed?

Reviewer #2 (Remarks to the Author):

In "Efficient sampling of high-dimensional free energy landscapes using adaptive reinforced dynamics", the authors presented improvements to their reinforced dynamics (RiD) enhanced-sampling approach and illustrated the benefits using a variety of systems. I appreciate the choice of the systems and their complexity, which go beyond the over-simplified alanine dipeptide test case that is still used by (too) many authors to benchmark new developments in the enhanced-sampling field. The presentation is also very clear and the approach will certainly be of interest to the computational chemists community. There are a couple of points that I would like the authors to address before I can recommend this paper for publication.

1) As the authors underlined, popular methods such as metadynamics (MetaD) and umbrella sampling are limited by the choice of 2/3 collective variables. However, there have been recent developments, especially as far as metadynamics is concerned, to alleviate this problem. As an example, Bias-Exchange MetaD and Parallel-Bias MetaD (PBMetaD) now enable the use of a large number of CVs in the MetaD scheme. In particular, PBMetaD does not require the use of a replica-exchange scheme as Bias-Exchange MetaD. I invite the authors to:
- acknowledge the existence of modern variations of metadynamics that address the problem of using a large set of CVs
- benchmark the efficiency of their proposed method against Bias-Exchange MetaD and/or PBMetaD using the peptoid and chignolin test cases with 3 and 18 CVs, respectively. This comparison can be performed by monitoring the deviation of the reconstructed free-energy from the reference (REMD) as a function of simulation time for the chosen methods. Both Bias-Exchange MetaD and PBMetaD are implemented in the PLUMED software, which the authors used for this study. Therefore, the very same CVs defined for their RiD simulations can be used in combination with the other enhanced-sampling techniques implemented in PLUMED

2) I invite the authors to report in Supplementary Informations the time series of the peptoid dihedral

angles (Fig. S2/S3) calculated on the continuous replicas that span different temperatures. As a matter of fact, plotting the reconstructed free energy (at 300K or any other temperature) as a function of simulation time is not sufficient in REMD to assess convergence. A "stationary" profile can indeed be observed even when the continuous replicas are not ergodic (for example are stuck in different regions of the conformational space) as a consequence of the exchanges.

Reviewer #3 (Remarks to the Author):

As the authors state in the abstract, the purpose of this paper is to add a few technical improvements to the reinforced dynamics of Ref. 7 in order to enable the method to generate free energy surfaces of very high dimension. The method targets collective variables (CVs), and the applications do, indeed, use a large number of these, in particular, over 100 CVs in the protein structure refinement problem, although in that example, the system never wanders too far from the native structure, even when it reaches the non-native 0.43 nm structure, due to the application of restraints. There is one school of thought that generating high-dimensional surfaces, such as those targeted here, has some value in that they can be subsequently refined to try to find essential pathways between conformational basins. The results are impressive, and the paper will likely be publishable. However, there are some issues the authors need to address before it can be reconsidered for publication.

1. There really is quite a large overlap between the method presented in this work and that of Ref. 7. I'm still uncertain what the new elements are (perhaps the use of the KL divergence?). The reader should not have to work this hard to figure out what the technical improvements are. This needs to be made very clear by the authors. Final determination of whether this paper rises to the level of Nature Comp. Sci. can only be made when it is clear what degree of new innovation is in this paper.

2. While I very much like how the authors use the neural network to represent the free energy surface in a way that can be used as a biasing technique, it is important for them to acknowledge that they were not the first ones to use neural networks to represent free energy surface, nor is theirs the only technique capable of handling large numbers of CVs. Ref. 6 of the present manuscript seems to be the first among those cited by the authors (Refs. 4-12) to employ neural networks in conjunction with enhanced sampling to represent and deploy free energy surfaces for computing observables. The NN in panel c of figure one closely resembles Figure 1 of Ref. 6. In addition, examples of the use of the methodology of Ref. 16 to explore and generate high-dimensional free energy surfaces [see Abrams and Vanden-Eijnden, PNAS 107, 4961 (2010), Chen et al. PNAS 112, 3235 (2015), Cendagorta et al. J. Phys. Chem. B 124, 3647 (2020), which also compared the performance of different machine learning models]. The authors need to do a better job of acknowledging the work that preceded theirs.

3. There are a number of things that I found unclear in the authors' presentation:

i. On page 4, the authors state the mean forces are evaluated by restrained MD simulations, and the authors cite Ref. 16 here. Ref. 16 is NOT a technique for performing restrained simulations. Ref. 16 is the temperature-accelerated MD (TAMD) approach, and while it involves a harmonic coupling between CVs and the s values, s is a dynamical variable. Are the simulations performed using TAMD or are they performed by actual restraints? With restraints, only local information about the free energy surface would be generated, and some scheme for selecting different values of s would be needed. On a high-dimensional landscape, this would seem to be essentially impossible unless the navigation scheme introduced previously by E and coworkers [J. Chem. Phys. 140, 164109 (2014)] is used. Otherwise, if TAMD is used, then authors need to say this, and they should also acknowledge the work of Abrams and Tuckerman J. Phys. Chem. B 112, 15742 (2008).

ii. On page 5, the authors describe a clustering strategy "that ensures that the CVs are selected to optimally represent the part of the CV space with large uncertainty." Does this mean refinement of the CVs, themselves, or is this just a process of selecting a subset of a priori selected CVs to help reduce the uncertainty in some part of the CV space? How often is this selection made in the course of a simulation?

iii. As a followup to ii, there are now numerous methods that employ machine learning to help identify optimal low-dimensional CV sets for characterizing different processes. Could the authors comment on the relative merits of their approach to these? After all, it seems that their approach requires a priori selection of CVs, which remains a significant ongoing challenge. I suspect that the neural networks used here could be used for further analysis to find a non-redundant low-dimensional CV space.

iv. The authors tout their use of deep neural networks to represent the free energy surface. However, it seems that they are really only using feed-forward networks with 4 layers. I would not characterize these as deep neural networks, and suspect the machine learning community would not either. I think these are characterized, rather, as shallow networks.

4. The CVs used here are all of the same type, i.e., backbone dihedral angles, which only characterize local conformational changes. Could the authors' approach be used for different types of systems, besides peptides and peptoids, e.g., materials systems with very different types of CVs or even biomolecules but with different types of CVs (angles, RMSDs, radii of gyration, native contacts, etc.)? That is, how much does the approach depend on the chosen CVs being as homogeneous as they are here?

Author Rebuttal, second revision:

Dear Dr. Pan,

We thank the reviewers for their time and effort taken to review this paper, and we appreciate their constructive comments. We also thank you for summarizing these comments. We have made significant efforts for addressing the reviewers' points, and we are confident that the revised manuscript will meet the publication standard of Nature Computational Science. In the following, we first address the points that are required to be substantially worked on. Then we reply in detail to all the issues raised by the three referees. Primary changes to the main manuscript are highlighted in blue.

Best regards,
Linfeng Zhang on behalf of all authors

Points to be substantially worked on

• As suggested by referee 2 and 3, please provide an improved discussion about novelty/technical improvements when compared with the most recent developments in the fields, e.g., other machine learning approaches, metadynamics methods, and your recent work on the topic [The Journal of chemical physics, 148(12):124113, 2018] We have made a more comprehensive description of the machine learning based enhance sampling approaches and other enhance sampling methods (page 3 lines 49-72 and page 4 lines 81-89, 94-97), and discussed the difference between the adaptive RiD and existing approaches, especially the original RiD method (see page 4 lines 101-109 and page 5 lines 116-122, 125-130). The novelty of the adaptive RiD compared to the original RiD lies in two aspects: (1) A clustering algorithm is applied to the configurations selected for labeling to reduce the number of configurations needed to represent the unexplored configuration space. (2) The uncertainty indicator and the bias are adaptively and automatically tuned by using the same clustering argument that quantifies the diversity of the explored configurations. These important improvements made possible extensive exploration of the free energy landscape of more than 100 collective variables. As required by the editor and the referee 3, we have explicitly stated the novelty on page 5 lines 125-130. See also our response to referees R2.1, R3.1 and R3.2.

• Referee 2 suggested that a better benchmark/comparison study in efficiency is needed. Specifically, this referee suggested benchmarking the efficiency against Bias-Exchange MetaD and/or PBMetaD using the peptoid and chignolin test cases with 3 and 18 CVs, respectively. We have banchmarked the efficiency of adaptive RiD against BEMetaD and PBMetaD in the cases of peptoid trimer and chignolin. In summary, the adaptive RiD outperforms BEMetaD and is comparable with PBMetaD. We reported the detailed comparison on page 11 lines 250-258 for peptoid, and on page 13 lines 273-284 for chignolin. See also the response R2.1. 1

• Referee 1 suggested to be more detailed about the conversion of a multi-dimensional potential energy surface to low dimensional PES (this will be important to predict variables to be compared with

experimental observations). We have provided detail description on the conversion from a high-dimensional FES to a low-dimensional one in the Supplementary information page 2. Please also see the response to referee R1.1.

• Referees have concerns about the current presentation of technical points. Please clarify those points as suggested, such as a better discussion about your deep neural network (referees 1 and 3 have concerns about the deepness of your framework), the details about whether TAMD is used, etc. We have clarified why our neural networks is in the category of deep neural network in the response R3.3 iv. We have made clear that we use restrained MD simulations to compute training labels and corrected our mistake in the citation (page 6 lines 145-148). Please also see our response R3.3 i.

• Jargons, like GDT-HA score as used in the protein structure prediction and refinement field, need to be better defined for our general audience. Please define jargons or specific scientific terms appropriately in your manuscript. We have added the full name of GDT-HA in the abstract and the following sentence on page 14 lines 296-297: “GDT-HA scores range from 0 to 100 where a higher score means a higher similarity between two protein structures.” Please also see our response R1.2.

• Please better discuss the generality and the limitations of your method. For example, referee 3 has concerns about points that how general your approach is for the treatment of different types of CVs. We explain the generality of adaptive RiD and our major aims in response R3.4. Adaptive RiD can handle different types of CVs and be employed in other situations, e.g. material systems, without substantial difficulty.

• Please provide better justification on your neural network structure; for instance, was an ablation study performed? How researchers from other fields can learn from your neural network model? These detailed information will be important to our multidisciplinary readership and also increase the potential impact of your paper. In the hyper-parameter tuning procedure, we first estimate the statistical error introduced by the restrained MD simulations. This error defines the highest accuracy achievable by our DNN representation. Then the hyper-parameters like the batch size, the start learning rate, and the learning rate decay speed, are tuned to minimize the number of epochs needed for achieving the optimal accuracy. In the first few adaptive RiD iterations, the width of the DNN may be chosen to be a relatively small value, for example, 100 each hidden layer. As the adaptive RiD goes on, more parts of the FES are explored and the training data accumulate, the accuracy of the DNN model is observed to decrease. This indicates that the current DNN architecture is not powerful enough compared with the complexity of the explored FES. We stop the adaptive RiD manually when the relative error reaches ~0.5, enlarge the DNN with a sub-network 2 initialized by the original DNN, and then retrain the new DNN. This can substantially reduce the error of DNN and the adaptive RiD can continue. It is noted that the strategy of gradually enlarging DNN helps to determine the size of DNN but is not necessary, because using a large enough DNN at the beginning would not cause any difficulty in

training the DNN. We have added the guideline for tuning the hyperparameters of the adaptive RiD in the Supplementary Information, and added a sentence on page 19 lines 408-409 referring to this guideline.

Response to Reviewer 1

Comment R1.1 It is not clear how low-dimensional free energy surface can be obtained from high dimensional one. As far as I understand the method, the output is a high-dimensional FES. For practical applications, such as prediction of protein thermostability or binding free energy of a protein-ligand complex, it is useful to have a 1D or 2D FES, not e.g. 20D. This is because we want to predict some variables that can be compared with an experiment. Conversion of 20D FES to 1D FES may be a nontrivial task. At the beginning I suspected this to be a weak point of the method and I suspected that this was the reason why authors focused on protein structure refinement, where this problem is less important. Later I found Fig. 2 where 1D FESes are presented, so the method can predict low dimensional FESes. I would ask authors to give more details on the conversion of a multidimensional FES to a low dimensional one.

Response to R1.1 We confirm that the output of RiD is a high-dimensional FES, and now provide the details on how a high-dimensional FES is converted to a low-dimensional one. We notice that the FES is a logarithm of the probability distribution, i.e.

$$A(\boldsymbol{s}) = -k_b T \ln p(\boldsymbol{s}), \tag{1}$$

thus the low-dimensional FES is computed from the marginal distribution of the highdimensional probability distribution corresponding to the high-dimensional FES. In this work, we use the Markov chain Monte Carlo (MC) method to calculate the marginal distribution. For example, we have an M-dimensional CV space s = (s1, . . . , sM), and want to calculate the FES A(s1) from A(s1, . . . , sM). Since we have the definition for the marginal distribution on s1

$$p(s_1) = \int p(s_1, \ldots, s_M) ds_2 \ldots ds_M, \tag{2}$$

the FES on s1 is given by

$$A(s_1) = -k_b T \ln \int e^{-\frac{1}{k_b T} A(s_1, \ldots, s_M)} ds_2 \ldots ds_M. \tag{3}$$

The integration in (2) or (3) is computed by Markov chain Monte Carlo (MC). Here, we carried out 2000 independent MC samplers and each lasts 106 steps. The example can be 3 easily generalized to any FES defined on a low-subspace of the high-dimensional CV space. In the revised manuscript, we added the above sentences to the Supplementary Information page 2.

Comment R1.2 The article is likely to be read mostly by molecular simulation experts. The authors use GDT-HA score without any explanation, even without explaining the abbreviation. Authors should add one or two sentences introducing GDT-HA to nonexpert in protein structure predictions. This should be addressed also in the abstract. The abstract should be self explanatory.

Response to R1.2 The global distance test (GDT), also written as GDT-TS to represent ”total score” or as GDT-HA to represent ”high accuracy”, is a measure of similarity between two protein structures, which is commonly used to compare the structures in protein prediction and refinement. GDT scores range from 0 to 100 where a higher score means a higher similarity. Now, GDT scores have been used as one of the standard measures in Critical Assessment of Structure Prediction (CASP). In the revised manuscript, we added the full name of GDT-HA in the abstract and the following sentence: “GDT-HA scores range from 0 to 100 where a higher score means a higher similarity between two protein structures.” on page 14 lines 296-297.

Comment R1.3 The switching function sigma is not clear to me. I understood that it is sigma(epsilon) = 0 for epsilon < e0, sigma(epsilon) = 1 for epsilon > e1, and sigma(epsilon) = (epsilon-e0) / (e1-e0) elsewhere. Is this correct? I would appreciate more clear explanation.
Response to R1.3 We use the switching function defined as:

$$\sigma(\epsilon) = \begin{cases} 1, & \epsilon < \epsilon_0 \\ \frac{1}{2} + \frac{1}{2}\cos\left(\pi \frac{\epsilon - \epsilon_0}{\epsilon_1 - \epsilon_0}\right), & \epsilon_0 < \epsilon < \epsilon_1 \\ 0, & \epsilon > \epsilon_1 \end{cases}$$

In the revised manuscript, we added this definition in equation 2 on page 8.

Comment R1.4 Evolution of e0 and e1 is quite confusing in Figure 1 and the related text. I would advice authors to use some symbol for iteration number, e.g. “t” (of course authors may use other symbol or notation) to replace e0=e0 by e0(t) = e0(0) and to replace e0 = e0 * 1.5 by e0(t) = e0(t-1) * 1.5 (maybe I’m wrong with the use of t, t-1 and 0, but this highlights the need for more clear mathematical description).

Response to R1.4 We thank the reviewer for this great suggestion. We renewed the Fig. 1 by replacing e0=e0 by e0(t) = e0(0) and e0 = e0 * 1.5 by e0(t) = e0(t-1) * 1.5.

Comment R1.5 I appreciate that authors provide the code at GitHub. As far as I can judge it is well established. I advise authors to also submit their input files via Plumed Nest. I understand that there

might be some obstacles because Plumed has been modified for this work. Regarding the modification, would it be possible to add it to official Plumed?

Response to R1.5 We have contacted the Plumed developers and found that there are still technically difficulties to merge our modification to the official Plumed. The main obstacle is that RiD links to the C++ interface of Tensorflow for efficient DNN operations, and the compiling of the C++ interface needs a lot of hacks thus is not easy to be made fully automatic. We are still working on the issue with Plumed developers and trying to find a solution.

We agree with the reviewer that the input files should be published via Plumed Nest. We have submitted our input files and the project ID is plumID:21.034. We have added the availability of the input files to the "Data and code availability" section on Page 23 lines 512-513.

Response to Reviewer 2

Comment R2.1 As the authors underlined, popular methods such as metadynamics (MetaD) and umbrella sampling are limited by the choice of 2/3 collective variables. However, there have been recent developments, especially as far as metadynamics is concerned, to alleviate this problem. As an example, Bias-Exchange MetaD and Parallel-Bias MetaD (PBMetaD) now enable the use of a large number of CVs in the MetaD scheme. In particular, PBMetaD does not require the use of a replica-exchange scheme as Bias-Exchange MetaD. I invite the authors to:
- acknowledge the existence of modern variations of metadynamics that address the problem of using a large set of CVs
- benchmark the efficiency of their proposed method against Bias-Exchange MetaD and/or PBMetaD using the peptoid and chignolin test cases with 3 and 18 CVs, respectively. This comparison can be performed by monitoring the deviation of the reconstructed free-energy from the reference (REMD) as a function of simulation time for the chosen methods. Both Bias-Exchange MetaD and PBMetaD are implemented in the PLUMED software, which the authors used for this study. Therefore, the very same CVs defined for their RiD simulations can be used in combination with the other enhanced-sampling techniques implemented in PLUMED

Response to R2.1 We regret our ignorance of the recently developed variances of the metadynamics. We have added the BEMetaD and PEMetaD methods in the Introduction and pointed out their ability of exploration in high-dimensional CV spaces on page 3 lines 58-72.

We benchmarked the BEMetaD and PBMetaD using the peptoid trimer and chignolin with 9 and 18 CVs, respectively. For BEMetaD, three parameter settings, denoted by BE0.2, BE0.5 and BE0.8, are considered. For PBMetaD one parameter setting (denoted by PB0.5) is investigated. We describe in detail the parameter settings in the Methods section of the text, see page 21-22 lines 456-471, and page 23 lines 491-493.

For peptoid trimer, adaptive RiD shows comparable efficiency with both BEMetaD and PBMetaD, in the sense that the free energy curves are all converged in about 1440 ns. For a quantitative comparison of the accuracy of different methods, an additional REMD of 300- 680 K and 400 ns each replica (4 times longer than other REMD simulations) is provided 5 as reference. We conclude that the accuracy of the adaptive RiD is comparable with the BEMetaD and PBMetaD (see Fig. S6 and Tab. SIII). In the revised manuscript, we discuss these results on page 11 lines 250-258.

For chignolin, we count the number of folding and unfolding events per microsecond and calculate the rate of the folding and unfolding transitions. The folding rate of all the method are comparable. The adaptive RiD presents a folding rate of 4.30 μs^{-1} . The highest rate is achieved by PBMetaD with 5.56 μs^{-1} , while the lowest rate is 2.31 μs^{-1} of BEMetaD (BE0.5). The unfolding rate of RiD is 4.30 μs^{-1} and is equal to its folding rate. This implies that the adaptive RiD successfully biased the system out of the global minimum (native state), and strongly supports the argument that the adaptive tuned bias encourages the system escape from deep minima of the FES. By contrast, the unfolding rates of BEMetaD degenerate to less than 1 μs^{-1} , which means that the system is trapped by the mimina of FES. One may argue that BEMetaD has already reached the native state and it is not necessary to sample the rest of the configuration space. This is true for Chignolin, a mini-protein with a relatively simple FES. For larger proteins, the depth of the local minima of FES may be the same or even deeper than the global mimimum of the chignolin FES, for example, a β-strand domain forms on the backbone of the protein, then the BEMetaD is likely to be trapped for a long time and have little chance of reaching the global mimimum. PBMetaD presents a satisfactory unfolding rate, 3.24 μs^{-1} , which is slightly lower than its folding rate. The detailed folding and unfolding rates are reported in Supplementary Information Tab. SV. In the revised manuscript, we show these results on page 13 lines 269-270 and 273-284.

Comment R2.2 I invite the authors to report in Supplementary Informations the time series of the peptoid dihedral angles (Fig. S2/S3) calculated on the continuous replicas that span different temperatures. As a matter of fact, plotting the reconstructed free energy (at 300K or any other temperature) as a function of simulation time is not sufficient in REMD to assess convergence. A ”stationary” profile can indeed be observed even when the continuous replicas are not ergodic (for example are stuck in different regions of the conformational space) as a consequence of the exchanges.

Response to R2.2 We thank the reviewer for reminding us of this point. The time series of the peptoid trimer dihedral angles (Fig. S2/S3) calculated on the continuous replicas that span different temperatures are shown in Fig. S4. In addition, the total number of transitions of six torsion angles ω1, ω2, ω3, ϕ1, ϕ2, ϕ3 of (s1pe)3 in different REMD simulations are shown in Tab. SII. It shows that there are tens of transitions in the REMD simulations of 300-430K and much more transitions are seen in the REMD simulations of 300-680K. It indicates that higher temperatures make it easier to converge. In the revised manuscript, we added the above sentences on page 11 lines 236-242.

Response to Reviewer 3

Comment R3.1 There really is quite a large overlap between the method presented in this work and that of Ref. 7. I'm still uncertain what the new elements are (perhaps the use 6 of the KL divergence?). The reader should not have to work this hard to figure out what the technical improvements are. This needs to be made very clear by the authors. Final determination of whether this paper rises to the level of Nature Comp. Sci. can only be made when it is clear what degree of new innovation is in this paper.

Response to R3.1 The major objective of the methodology developed in this work is to allow using neural network represented FES to facilitate the exploration of truly highdimensional free energy landscapes. This is in principle possible for many existing works, but in practice, to our best knowledge, molecular dynamics simulations using more than 100 CVs have not been reported. Original RiD has been shown to be successful in exploring spaces with no more than 20 CVs, but its efficiency deteriorates quickly with more CVs: It has been observed the RiD often gets trapped in the deep local minima for higher dimensional systems. This difficulty is caused by the following two factors: (1) The probability of visiting the neighborhood of a local minimum in a high-dimensional CV space is much lower than that in lower dimensional examples. Thus a random batch of RiD samples is not enough for reconstructing the landscape near the local minimum of a high-dimensional FES. (2) The biasing mechanism is too rigid for exploration, and particularly for escaping from deep local minima. It was the novelty of the adaptive RiD method that makes our goal a reality. The novelty lies in two aspects: (1) The introduction of the clustering algorithm in the selection of the training data. All the explored data with a large indicated error are clustered by an agglomerative cluster algorithm, and only one configuration is randomly selected from each cluster for labeling as the training data. This makes the training dataset distributed in an optimal way in the un-explored configuration space. (2) The error indicator and the bias are adaptively tuned by using the clustering algorithm. We quantify the diversity of the explored configurations by the number of clusters on the exploration trajectory. A low diversity usually means that the system is trapped in a deep local minima of the FES, so the uncertainty levels are increased to encourage the escaping from the local minima. It is noted that the adjustment of uncertainty levels is fully automatic, thus the adaptive RiD only needs little human intervention and no a priori knowledge on the energy barriers of the FES.

To show how the adaptive RiD method overcomes the intrinsic difficulties faced by the original RiD method, we compare their performances in the Methods section and in Fig. 6. In detail, the original RiD has a fixed uncertainty level interval. It is easy for RiD to be stuck in a FES local minima surrounding by high barriers, and the generation of new data will be quite slow in that case, leading to an inefficient biasing. When this happens, adaptive RiD will elevate uncertainty levels to enable a large phase space to be biased, which help climbing over the barriers at the cost of lowering the accuracy of bias potential. When the system escapes from the local minima, the uncertainty levels will be reset to ensure an

accurate description of the FES. In the original RiD method, effort is needed for tuning the uncertainty levels to achieve a balance between the accuracy of FES and the efficiency of exploration. The adaptive RiD, on the other hand, makes the parameter tuning automatic and adapted to the FES, thus it minimizes the expertise needed to use the method. According to our results in Figure 6, one can easily check that simulations under adaptive RiD schemes have more frequent transitions and every transition comes with an elevation of uncertainty levels (panel d and g). It is this adaptive adjustment that makes RiD efficient in new fields like the folding of complex proteins and protein structure refinement.

To make the contribution of adaptive RiD clear, we have commented on the drawback of the original RiD on page 4 lines 101-109 of the revised manuscript, and explicitly state the novelty of the adaptive RiD on page 5 lines 125-130.

Comment R3.2 While I very much like how the authors use the neural network to represent the free energy surface in a way that can be used as a biasing technique, it is important for them to acknowledge that they were not the first ones to use neural networks to represent free energy surface, nor is theirs the only technique capable of handling large numbers of CVs. Ref. 6 of the present manuscript seems to be the first among those cited by the authors (Refs. 4-12) to employ neural networks in conjunction with enhanced sampling to represent and deploy free energy surfaces for computing observables. The NN in panel c of figure one closely resembles Figure 1 of Ref. 6. In addition, examples of the use of the methodology of Ref. 16 to explore and generate high-dimensional free energy surfaces [see Abrams and Vanden-Eijnden, PNAS 107, 4961 (2010), Chen et al. PNAS 112, 3235 (2015), Cendagorta et al. J. Phys. Chem. B 124, 3647 (2020), which also compared the performance of different machine learning models]. The authors need to do a better job of acknowledging the work that preceded theirs.

Response to R3.2 We are sorry for not acknowledging the pioneering works completely. We have updated our statements and references in the Introduction section. We have explicitly pointed out that we are not the first to use DNN to enhance the sampling of the system, nor the first method to handle high-dimensional CV space. The revised parts of the manuscript are listed as follows: page 3 lines 49-72, page 4 lines 81-89 and lines 94-97.

Comment R3.3 There are a number of things that I found unclear in the authors' presentation:

i. On page 4, the authors state the mean forces are evaluated by restrained MD simulations, and the authors cite Ref. 16 here. Ref. 16 is NOT a technique for performing restrained simulations. Ref. 16 is the temperature-accelerated MD (TAMD) approach, and while it involves a harmonic coupling between CVs and the s values, s is a dynamical variable. Are the simulations performed using TAMD or are they performed by actual restraints? With restraints, only local information about the free energy surface would be generated, and some scheme for selecting different values of s would be needed. On a high-dimensional landscape, this would seem to be essentially impossible unless the navigation scheme

introduced previously by E and coworkers [J. Chem. Phys. 140, 164109 (2014)] is used. Otherwise, if TAMD is used, then authors need to say this, and they should also acknowledge the work of Abrams and Tuckerman J. Phys. Chem. B 112, 15742 (2008).
ii. On page 5, the authors describe a clustering strategy ”that ensures that the CVs are selected to optimally represent the part of the CV space with large uncertainty.” Does this mean refinement of the CVs, themselves, or is this just a process of selecting a subset of a priori selected CVs to help reduce the uncertainty in some part of the CV space? How often is this selection made in the course of a simulation?
iii. As a follow up to ii, there are now numerous methods that employ machine learning to help identify optimal low-dimensional CV sets for characterizing different processes. Could the authors comment on the relative merits of their approach to these? After all, it seems that their approach requires a priori selection of CVs, which remains a significant ongoing challenge. I suspect that the neural networks used here could be used for further analysis to 8 find a non-redundant low-dimensional CV space.
iv. The authors tout their use of deep neural networks to represent the free energy surface. However, it seems that they are really only using feed-forward networks with 4 layers. I would not characterize these as deep neural networks, and suspect the machine learning community would not either. I think these are characterized, rather, as shallow networks

Response to R3.3
i. We are sorry for the confusing statement on the restrained MD simulation as well as the wrong reference. Restrained MD simulations is used in RiD as a tool for labeling. What we actually do is to add a harmonic potential between the target CV value and the instantaneously CV, and the gradient of the FES at the target CV is estimated by the mean restraining force. The reference Ref.16 was not cited properly, it should have been ”Single-sweep methods for free energy calculations” by the same authors published in 2008, who used the restrained MD to estimate the gradients of an FES. We have revised the manuscript on page 6 lines 145-147.
ii. We are sorry for the ambiguity in the description of the CV proposal-selection procedure. The CVs proposal-selection procedure is performed in every adaptive RiD iteration after the exploration step. In detail, we calculate the variance of the outputs of the DNN models on each explored configuration. The configurations with high variance are proposed for selection, because it implies a deficient training of the DNN models on that configuration. Then the proposed configurations are selected by the clustering argument, and finally the selected configurations are sent to the labeling step. We have updated the manuscript on page 8 lines 165-169 to make this point more clear.
iii. We acknowledge the reviewer for the valuable comment. There is a great volume of work on the identifying a low-dimensional CV space from high-dimensional MD simulation trajectories. In this work we did not try to solve the problem on how to design an optimal set of CVs for the enhanced sampling problems via machine learning, but to tackle the problem of efficient sampling and explore the relevant conformations given a predefined set of CVs. As demonstrated in the manuscript, adaptive RiD is a competitive method, especially for the sampling problems in high-dimensional CV space. This feature

makes it less sensitive the the choice of CVs, because the candidate CV can be all included to form a high-dimensional CV space, and adding a non-relevant CV would not degenerate the performance of RiD, as has been discussed in the original RiD paper (Zhang et.al JCP 148, 124113 (2018)).
iv. We appreciate the reviewer for questioning this concept, which could confuse most of people by its appearance. We want to clarify that deep neural networks (DNN) refers to the artificial neural network architectures using multiple (≥ 2) hidden layers between the input and output layers, which is opposite to the concept of the shallow neural network that owns only one hidden layer. This concept is widely accepted in literature, and the differences between the deep and shallow neural networks have been theoretically investigated. For example, Delalleau and Bengio [NeurIPS 24, 666–674 (2011)] pointed out that there exist families of functions that can be represented much more efficiently with a deep network than with a shallow one. Therefore, the expressive power of neural networks with more than one hidden layer is qualitatively different than those with one hidden layer. There are a lot of follow-up work demonstrating the benefit of using deep neural networks, for example, Liang and Srikant, ICLR (2017), Telgarsky, Conference on learning theory (2016), Montanellia, Yang and Du arXiv:1903.00735 (2019). As an another example, Restricted Boltzmann Machine (RBM) and Deep Boltzmann Machine (DBM), which are important 9 extraction methods, have only 2 or 3 layers in total. The NNs employed in RiD, with 4 hidden layers, are in the scope of deep neural network.

Comment R3.4
The CVs used here are all of the same type, i.e., backbone dihedral angles, which only characterize local conformational changes. Could the authors' approach be used for different types of systems, besides peptides and peptoids, e.g., materials systems with very different types of CVs or even biomolecules but with different types of CVs (angles, RMSDs, radii of gyration, native contacts, etc.)? That is, how much does the approach depend on the chosen CVs being as homogeneous as they are here? Response to R3.4 Actually, we tested other types of CVs on chignolin before, such as RMSDs, radii of gyration, hydrophobic contacts, alpha helicity, beta similarity, dihedral correlations, hydrogen bonds, etc. In addition, we also tested the CVs of distances combined with dihedral angles in peptide binding. All showed that RiD works well. In this paper, our major interest is on the huge number of CVs, e.g. tens or hundreds, so we just choose all the dihedral angles as CVs straightforwardly. RiD can also be employed in material systems without substantial difficulty. As the reviewer mentioned, dihedral angles can only reflect local structure information, but when employing all dihedral angles of the system, NN can also extract information of global structure changes.

Decision letter, third revision:

Dear Dr. Zhang,

Thank you for submitting your revised manuscript "Efficient sampling of high-dimensional free energy

landscapes using adaptive reinforced dynamics" (NATCOMPUTSCI-21-0325C). It has now been seen by the original referees and their comments are below. The reviewers find that the paper has improved in revision, and therefore we'll be happy in principle to publish it in Nature Computational Science, pending minor revisions to satisfy the referees' final requests and to comply with our editorial and formatting guidelines.

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

TRANSPARENT PEER REVIEW
Nature Computational Science offers a transparent peer review option for new original research manuscripts submitted from 17th February 2021. We encourage increased transparency in peer review by publishing the reviewer comments, author rebuttal letters and editorial decision letters if the authors agree. Such peer review material is made available as a supplementary peer review file. Please state in the cover letter ‘I wish to participate in transparent peer review’ if you want to opt in, or ‘I do not wish to participate in transparent peer review’ if you don’t. Failure to state your preference will result in delays in accepting your manuscript for publication.
Please note: we allow redactions to authors’ rebuttal and reviewer comments in the interest of confidentiality. If you are concerned about the release of confidential data, please let us know specifically what information you would like to have removed. Please note that we cannot incorporate redactions for any other reasons. Reviewer names will be published in the peer review files if the reviewer signed the comments to authors, or if reviewers explicitly agree to release their name. For more information, please refer to our FAQ page.

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

Sincerely,

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

ORCID
IMPORTANT: Non-corresponding authors do not have to link their ORCIDs but are encouraged to do so. Please note that it will not be possible to add/modify ORCIDs at proof. Thus, please let your co-authors know that if they wish to have their ORCID added to the paper they must follow the procedure

described in the following link prior to acceptance:
https://www.springernature.com/gp/researchers/orcid/orcid-for-nature-research

Reviewer #1 (Remarks to the Author):

The authors significantly improved the manuscript and addressed all comments I made. I support its publication.

Reviewer #2 (Remarks to the Author):

The authors performed a substantial revision of the paper and satisfactorily addressed my previous concerns. I recommend publication of the manuscript in the current form.

Reviewer #3 (Remarks to the Author):

I appreciate all of the changes made by the authors. They have addressed the comments in my original review. The only remaining change that needs to be made to the newly added text is the following: On page 3, the description of the TAMD method is not correct. In fact, what they are describing is the original temperature-accelerated method (adiabatic free energy dynamics) of Rosso et al. J. Chem. Phys. 116, 4389 (2002), which should be cited by the authors. However, in the method of Refs. 4 and 5, what is done is that the CVs are coupled harmonically to a set of extended phase-space variables, in fact, the s variables in the first paragraph of Sec. II, and it is the s variables that are coupled to a high temperature and adiabatic decoupling. The authors will need to fix this description before the paper can be accepted.

Author rebuttal, third revision:

Dear Dr. Pan,

We thank the reviewers for their encouraging comments. In the following we provide a point-by-point response to the reviewers' comments. We also provide a document that highlights the changes that we have made.

Best regards,

Linfeng Zhang on behalf of all authors

Response to Reviewers 1 and 2

Response to R12.0 We thank the reviewers for their supportive comments.

Response to Reviewer 3

Comment R3.1 On page 3, the description of the TAMD method is not correct. In fact, what they are describing is the original temperature-accelerated method(adiabatic free energy dynamics) of Rosso et al. J. Chem. Phys. 116, 4389 (2002), which should be cited by the authors. However, in the method of Refs. 4 and 5, what is done is that the CVs are coupled harmonically to a set of extended phase-space variables, in fact, the s variables in the first paragraph of Sec. II, and it is the s variables that are coupled to a high temperature and adiabatic decoupling. The authors will need to fix this description before the paper can be accepted.

Response to R3.1 We have revised the manuscript according to the reviewer's suggestion. See the second paragraph of the Introduction part: "An important advance was the temperature accelerated MD ..."

Final Decision Letter:

Dear Dr Zhang,

We are pleased to inform you that your Article "Efficient sampling of high-dimensional free energy landscapes using adaptive reinforced dynamics" has now been accepted for publication in Nature Computational Science.

In approximately 10 business days you will receive an email with a link to choose the appropriate publishing options for your paper and our Author Services team will be in touch regarding any additional information that may be required.

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

Authors may need to take specific actions to achieve compliance with funder and institutional open access mandates. For submissions from January 2021, if your research is supported by a funder that requires immediate open access (e.g. according to Plan S principles)

then you should select the gold OA route, and we will direct you to the compliant route where possible. For authors selecting the subscription publication route our standard licensing terms will need to be accepted, including our self-archiving policies. Those standard licensing terms will supersede any other terms that the author or any third party may assert apply to any version of the manuscript.

You will not receive your proofs until the publishing agreement has been received through our system.

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

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

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

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

If you have queries at any point during the production process then please contact the production team at rjsproduction@springernature.com. Once your paper has been scheduled for online publication, the Nature press office will be in touch to confirm the details.

Content is published online weekly on Mondays and Thursdays, and the embargo is set at 16:00 London time (GMT)/11:00 am US Eastern time (EST) on the day of publication. If you need to know the exact publication date or when the news embargo will be lifted, please contact our press office after you have submitted your proof corrections. Now is the time to inform your Public Relations or Press Office about your paper, as they might be interested in promoting its publication. This will allow them time to prepare an accurate and satisfactory press release. Include your manuscript tracking number NATCOMPUTSCI-21-0325D and the name of the journal, which they will need when they contact our office.

About one week before your paper is published online, we shall be distributing a press release to news

organizations worldwide, which may include details of your work. We are happy for your institution or funding agency to prepare its own press release, but it must mention the embargo date and Nature Methods. Our Press Office will contact you closer to the time of publication, but if you or your Press Office have any inquiries in the meantime, please contact press@nature.com.

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

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

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

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

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

We look forward to publishing your paper.

Best regards,

Jie Pan, Ph.D.

Associate Editor
Nature Computational Science