nature portfolio

Peer Review File

Exploring Chemistry and Catalysis by Biasing Skewed Distributions via Deep Learning

Corresponding Author: Professor GiovanniMaria Piccini

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

Attachments originally included by the reviewers as part of their assessment can be found at the end of this file.

Version 0:

Reviewer comments:

Reviewer #1

(Remarks to the Author)
The manuscript presents Loxodynamics, a novel enhanced sampling scheme meant to explore chemical pathways in a data-driven way. At the core of the method is a deep autoencoder architecture (Skewencoder) featuring a bespoke loss function which uses the skewness of the latent space distribution to identify chemical reaction coordinates (CV). This CV is iteratively updated and used to bias the dynamics via a static half wall potential.
The problem of identifying chemical reaction pathways is extremely relevant and any development in this area is of interest to many researchers. The proposed method introduces some interesting new approaches, such as the clever combination of deep learning and statistical analysis of the anharmonicity to guide the sampling. I believe it deserves publication, especially after the following minor points have been addressed.

1. Since it is important for the method to properly work, it should be mentioned in the main paper that one has to have a way to identify the metastable states. At least in the SI, please provide an example of how the metastable states were identified in the systems presented.

2. It would be interesting to report, in the main text or in the SI, the trajectory of the learned CV, i.e. of the latent space. This would give a more intuitive idea of the anharmonicities the method is set up to follow. How does this CV compare with the ones shown in the plots? How much does this CV change at each iteration?

3. My understanding is that the skewness gamma of the latent space can switch sign between iterations, spontaneously breaking the symmetry of the loss L_skew. This leads to the annoyance of having to keep track of the sign of gamma when defining the biasing wall. But, instead of that, one could simply change the definition of the CV to ` s'_LS = sign(gamma_ls) * s_LS `, this would simplify eq. 8, fig. S1, and some of the code, since one would always need a lower wall.

4. Could you briefly comment on the choice of eq. 4 for the skewness loss? Why not a simple gaussian? Why not an asymmetric function that would always push the skewness of the latent space to be of the same sign?

5. Please be more explicit in the mathematical formulation on the loss function, eq. 2, and report it in full, e.g in the methods section. This would help the reader to understand. The loss L should be a function of the parameters of the NN, theta, and of the descriptors x. Gamma is computed on the latent space y, thus gamma(y) is also a function of the parameters theta.

6. “A larger harmonic constant was applied when steering the reaction in the direction [...] to account for the higher barrier” how do you choose the harmonic constant? Is it difficult?

7. It’s not clear how Loxodynamics finds multiple paths. Trial and error? When should I stop? Does this mean that the first path I find could be a higher FES pathway, and not the most common transition state?

(Remarks on code availability)

The code is clearly written and reasonably documented. It contains all the necessary tools to reproduce the results of the paper or use the proposed method on different systems, including a Jupyter notebook tutorial.

Reviewer #2

(Remarks to the Author)
The manuscript by Zhang and Piccini presents an original and promising approach to reaction discovery based on the skewness of probability distribution using an autoencoder framework. The idea behind the method—Loxodynamics—is clearly explained and introduces a creative use of probability distribution asymmetry to guide sampling. However, several methodological and validation aspects require clarification and improvement before the work can be recommended for publication.

Major Comments

>> Related Work and Literature.
The paper lacks a thorough discussion of prior literature on reaction discovery using machine learning collective variables, especially autoencoder-based methods. Relevant prior work includes:
- Chen & Ferguson (https://doi.org/10.1063/1.5023804) demonstrated the use of autoencoders for interleaved CV discovery and enhanced sampling, iteratively expanding the sampled space.
- Ketkaew & Luber (https://doi.org/10.1021/acs.jcim.2c00883) applied autoencoders to generate blind CVs for chemical reactions starting only from reactants
- Raucci et al. (https://doi.org/10.1021/acs.jpclett.1c03993) proposed a method for blind reaction discovery using general-purpose CVs with OPES explore.

>> Method Validation
- Comparison with standard autoencoders: what advantage does the “Skewencoder” offer over a standard autoencoder with the same architecture? A comparative experiment (e.g., on the toy model) would help clarify its contribution.
- Minimum Free Energy Pathways: How can it be guaranteed that the discovered paths correspond to minimum free energy paths? Except for the first toy model, the others only have one pathway. The last example has two pathways, and the fact that it can also find a higher pathway seems to contradict this aspect. Also, it seems that this will depend on the hyperparameters (e.g., \kappa, \epsilon)? Is prior knowledge of energy barriers required to set them?
- The discussion section claims a wide range of potential applications, but the validation is limited to a few relatively simple cases (except for the last one). A more realistic assessment of current limitations, along with suggestions for how to transition from exploration to detailed mechanistic modeling, would strengthen the manuscript. For instance: after discovery, can the CVs produced by Skewencoder be used in conventional biased simulations to compute free energy surfaces?

>> Methodological Details
- It is unclear whether \kappa is chosen by the user or adapted during the iterations (as suggested. In the SN2 example, it is said: “Due to the asymmetric energy barriers… a larger harmonic constant was applied … to account for the higher barrier.” Does setting this parameter require prior barrier knowledge?
- The rationale behind the inclusion of \sigma in the wall position (i.e., \mu + \sigma + \epsilon) is not explained in the main text. Figure 1 should also reflect this logic explicitly. Also, the role of \epsilon seems important for determining the degree of exploration, but is not discussed. A discussion of how to choose these parameters and what effect they have is needed.
- The claim that the reaction discovery is “unbiased” is misleading. Since the skewness of the biased distribution depends on the applied parameters of the harmonic bias and no reweighting is performed, the method introduces a directional sampling artifact.
- A complex procedure to decide whether to warm-start the autoencoder or not is put forward, but not justified. A comparison with an exploration training only from scratch on the toy model would clarify its impact. Also, this procedure requires a criterion to determine whether a new metastable state has been reached, which is not described for the examples reported. This contradicts the claim of minimal user intervention and should be detailed more clearly.
- Section A of the methods could benefit from a few revisions. I think that the workings of neural network layers are now part of common knowledge and so can be moved to the SI. On the other hand, it is not clear if the details about the L_{skew} are just conceptual considerations (as it seems to me looking at the code) or technical details that need to be implemented to improve the performance of the method. Whichever the case, I think this should be made clear to the reader.
- In general, it would be much better to include in the methods section a summary of the method, including the loss functions (which are only defined in the introduction) and the workflow of the iterative procedure (which is again discussed in the introduction and summarized only in the SI).
Below there are other comments specific to the applications presented.

>> Toy Model
- The evolution of the skewness and the mean of the CV across iterations is not shown. Including this would help illustrate how the method progresses and adapts the biasing direction based on the measured skewness?
- How does the exploration change as a function of the chosen hyperparameters, \kappa and \epsilon? Showing supplementary examples would help clarify how to set them.
- Although the potential includes two distinct paths, only one is sampled. Can the method recover both, depending on the hyperparameters used? This would also help clarify how this is obtained in the chabazite case.

>> SN2 Reaction
- The input CV space essentially consists of d_{Cl} and d_F, which already capture the transition without needing any

combination of them. In my opinion, this system should be moved to the SI, as it does not significantly validate the method.

>> Chabazite
- How is the second pathway found? In the SI it is said that different elastic constants are used to reflect barrier differences. Since barrier information is usually unknown, how would one proceed in general?
- In the stepwise mechanism, how is backward motion prevented? Are additional walls required? Would this become limiting for more complex catalytic cycles involving many intermediate steps?

>> Minor Comments
- Many figure captions are not self-contained. For example, the meaning of color schemes and symbols (e.g., Figure 1 symbols, Figure 3 colorbar, color of scatter plots, Figure 4 color of scatter plots, Figure S1 symbols, Figure S2 colorbars, Figure S4 color of scatter plots, colorbars) is only defined in the main text and should be reported in the captions.
- The provided code is well-structured and detailed, which ensures reproducibility. However, from the SI and the code, it appears to rely extensively on the mlcolvar library, which is not cited in the main text.
- When quantities are discussed in the main text, refer to the specific equation number to help guide the reader (e.g., L_{skew} in methods sec. A).

(Remarks on code availability)
The provided code is well-structured and detailed, which ensures reproducibility. However, from the SI and the code, it appears to rely extensively on the mlcolvar library, which is not cited in the main text.

Reviewer #3

(Remarks to the Author)
I co-reviewed this manuscript with one of the reviewers who provided the listed reports. This is part of the Nature Communications initiative to facilitate training in peer review and to provide appropriate recognition for Early Career Researchers who co-review manuscripts.

(Remarks on code availability)

Reviewer #4

(Remarks to the Author)
In this manuscript, Zhang and Piccini propose a novel approach to explore reactive paths using an encoder that includes the skewness of the probability distribution observed in a (set of) collective variables into its loss function. I find the manuscript to be clearly written and the method is well-explained. The authors, moreover, support their method with a handful of examples ranging from toy potentials to heterogeneous catalysis. While I would most certainly recommend publication in a more specialized journal, I believe that, in their present form, the results might not be sufficiently groundbreaking or of broad enough applicability given the multidisciplinary audience of Nature Communications. My main points of concern, in this regard, are the following:
1. The authors state “Moreover, upon algorithm termination, the accumulated samples provide sufficient data to construct a latent space that effectively approximates a meaningful CV”. While it looks like that is the case with the simple toy systems, it would be interesting to see if the conclusion holds when going to the more realistic reactions (e.g., plotting the isolines on the d6 vs d4 graph in Figure 4). I would find a reaction finder that also provides an (ok-ish) CV to sample the reaction with enhanced sampling quite interesting.
2. Concerning the last case study, I find it presently unclear what the difference was between the simulations that led to the exploration of the concerted vs stepwise mechanisms (aside from the wall kappa, as reported in the SI). Was sampling the two paths in the two simulations a coincidence? How different are they in energy according to the semiempirical method used by the authors? Also, in the stepwise mechanism, how likely it is to actually resample the backward reaction from the surface species to ethene instead of going forward?
3. I find that some comparison with literature is missing. For example, how does the method proposed by the authors compare with OPES-based exploration methods (e.g., https://pubs.acs.org/doi/10.1021/acs.jpclett.1c03993) in terms of cost and ability of finding new reaction products?
4. The reported case studies, while comprehensive, are all pretty simplistic. To remain in zeolite catalysis, I guess that using a larger adsorbed molecule (with, e.g., double bonds) could open up quite a lot more reactive paths (cyclization, cracking, and so on). Going back to comment 2, it would be interesting to see which paths is the method able to discover.
If the authors can reply to these comments, the manuscript might become suitable for Nature Communications. I additionally have the following minor comment:
1. I am not fully sure I understand all the passages behind warm-start training explained in the SI. Am I correct in stating that warm start is used unless the system falls into a new basin of attraction, in which case everything restarts from scratch? In the final part, the authors mention the case in which the system falls back to the original minimum. I do not immediately see how training from scratch would suddenly push the system towards another barrier.

(Remarks on code availability)

Version 1:

Reviewer comments:

Reviewer #1

(Remarks to the Author)
The authors addressed all my concerns and the new version of the paper is in excellent shape, especially the much improved SI. I believe it makes a significant contribution to the field and I recommend it for publication.

(Remarks on code availability)

Reviewer #2

(Remarks to the Author)
See attached file

(Remarks on code availability)

Reviewer #3

(Remarks to the Author)
I co-reviewed this manuscript with one of the reviewers who provided the listed reports. This is part of the Nature Communications initiative to facilitate training in peer review and to provide appropriate recognition for Early Career Researchers who co-review manuscripts.

(Remarks on code availability)

Reviewer #4

(Remarks to the Author)
I appreciate the effort that the authors have undertaken to reply to all the reviewers' comments. I find the additional examples and information included in the manuscript and SI to drastically improve the reader's understanding of the method as well as its relevance. Therefore, I would be happy to recommend the manuscript for publication in Nature Communications.
I only have one final remark for the authors: in section S3, a comparison between the free energy profiles obtained from a chemically-inspired CV and the CV from Loxodynamics is reported (Figure S12). However, the free energy profile is CV dependent and comparing two profiles directly is not per se meaningful (see, e.g., https://doi.org/10.1016/j.jcat.2020.04.015 or https://doi.org/10.1063/5.0083423). I would find it interesting to evaluate the actual CV-independent free energy difference between the two metastable states from the simulations results, to showcase that the Loxodynamics CV is (likely) fine and discrepancies in the free energy profile with chemically-inspired CVs are meaningless.

(Remarks on code availability)

Response to the reviewer's comments to manuscript **NCOMMS-25-28852-T**

"Exploring Chemistry and Catalysis by Biasing Skewed Distributions via Deep Learning"

Zhikun Zhang, GiovanniMaria Piccini

October 2, 2025

We would like to sincerely thank the reviewers for their thorough, thoughtful, and constructive criticism. The comments we received were both rigorous and fair, and we greatly appreciated the opportunity to carefully address each of them. Following your suggestions has been an intellectually rewarding process, and we are confident that the manuscript has improved substantially in clarity, scope, and scientific impact as a result. We are very glad to have incorporated all your requests, and we hope you will find the revised version significantly strengthened thanks to your valuable input.

Reviewer #1

Remarks to the Author:

The manuscript presents Loxodynamics, a novel enhanced sampling scheme meant to explore chemical pathways in a data-driven way. At the core of the method is a deep autoencoder architecture (Skewencoder) featuring a bespoke loss function which uses the skewness of the latent space distribution to identify chemical reaction coordinates (CV). This CV is iteratively updated and used to bias the dynamics via a static half wall potential.

The problem of identifying chemical reaction pathways is extremely relevant and any development in this area is of interest to many researchers. The proposed method introduces some interesting new approaches, such as the clever combination of deep learning and statistical analysis of the anharmonicity to guide the sampling. I believe it deserves publication, especially after the following minor points have been addressed.

1. Since it is important for the method to properly work, it should be mentioned in the main paper that one has to have a way to identify the metastable states. At least in the SI, please provide an example of how the metastable states were identified in the systems presented.

#Response#:

We thank the reviewer for this valuable comment and appreciate the suggestion. In response, we have added a new subsection "Metastable states identification" (Sec. S1. C, Page 5, Line 78, see SI) detailing the methodology used to identify metastable states in both the reaction system and the model system mentioned in our study. Additionally, we have explicitly stated the criterion in the main text (Page 11, Line 759-760, see main text).

2. It would be interesting to report, in the main text or in the SI, the trajectory of the learned CV, i.e. of the latent space. This would give a more intuitive idea of the anharmonicities the method is set up to follow. How does this CV compare with the ones shown in the plots? How much does this CV change at each iteration?

#Response#:

We thank the reviewer for this insightful comment. As suggested, we have added CV contour plots for representative simulation iterations in the model systems to the SI (Figure. S4, Page 8, see SI), illustrating how the latent space CV evolves over time. Furthermore, we report that the Loxodynamics biases enhance sampling toward saddle points, thereby capturing local anharmonic features more effectively.

3. My understanding is that the skewness gamma of the latent space can switch sign between iterations, spontaneously breaking the symmetry of the loss $\mathcal{L}_{skew}$. This leads to the annoyance of having to keep track of the sign of gamma when defining the biasing wall. But, instead of that, one could simply change the definition of the CV to $s'_{LS} = \mathbf{sign}(\gamma_{ls}) * s_{LS}$, this would simplify eq. 8, fig. S1, and some of the code, since one would always need a lower wall.

#Response#:

We deeply appreciate this comment from the reviewer and we embraced the valuable suggestion. In this regard, the simplified format proposed by the reviewer is mathematically equivalent to the formula presented in Equation 8 of the original manuscript. To further analyze these aspects, we conducted comparative experiments on three model PES to substantiate this assertion, substituting the initial wall application with a newly simplified version while consistently imposing a lower boundary condition. The results are entirely reproducible using this revised harmonic wall application format. Results can be found in either github code or cosine database of the corresponding manuscript. Consequently, we have amended the main text's formula concerning wall application according to the recommended form (Equation 7, Page 10, Line 687-688). Nonetheless, we intend to retain the original formulation within the algorithm description section of the SI. The terms "lower wall" and "upper wall" possess greater physical significance, enhancing comprehension of how skewness corresponds to local anharmonicity and how Loxodynamics bias is applied accordingly (Figure S1, Page 4, see SI). This approach may facilitate readers in recognizing coherence and improve their understanding.

4. Could you briefly comment on the choice of eq. 4 for the skewness loss? Why not a simple gaussian? Why not an asymmetric function that would always push the skewness of the latent space to be of the same sign?

#Response#:

The point raised by the reviewer is very important and we are willing to clarify it in our response here as well as by adding new information in the revised SI (vide infra).

Initially, we wish to clarify that the incorporation of the skewness term into the loss function aims to optimize a latent space such that projections of "local dataset" samples within this space exhibit relatively maximal skewness. Therefore, the skew loss term should serve as a metric for the magnitude (i.e., absolute value) of sample skewness projected onto the latent space, excluding its sign. Upon completing optimization based on this absolute value, the sign of optimized projected sample skewness reveals the slowest ascent direction within the latent space. As noted in Point 3 by Reviewer #1, with certain mathematical manipulations, one can flip this sign, implying how system progresses towards a reaction channel is reflected entirely by this absolute value of sample skewness. However, employing an asymmetric function that consistently forces latent space skewness to maintain a uniform sign affects loss function optimization and may result in unforeseen performance degradation during exploration.

Nonetheless, various symmetric forms of skewness-related loss functions can theoretically serve as valid alternatives to the current application within the Skewencoder framework. New tests and analyses have been conducted on different model PES for comparison, as detailed in the revised SI (Section S2.A.1, Page 8-9). We evaluated the performance of Skewencoder using three distinct formulations of the skewness loss term: 1) the original implementation; 2) a Gaussian-like form, represented as $\exp(-\gamma^2)$*; and 3) a simplified quadratic form of projected sample skewness, expressed as* $-\gamma^2$*. Considering criteria primarily focused on computational efficiency and training resource consumption, our original implementation of the skewness loss function demonstrated superior performance among all tested formats. This comparison substantiates our assertion that the initially designed loss function offers advantages in facilitating more efficient training and sampling processes.*

5. Please be more explicit in the mathematical formulation on the loss function, eq. 2, and report it in full, e.g in the methods section. This would help the reader to understand. The loss L should be a function of the parameters of the NN, theta, and of the descriptors x. γ **is computed on the latent space y, thus** $\gamma(y)$ **is also a function of the parameters theta.**

#Response#:

We appreciate the reviewer for highlighting this issue. In response, we have provided a more comprehensive explanation of the loss function. Specifically, the autoencoder reconstruction loss $\mathcal{L}_{Autoencoder}$ *is dependent on the global input dataset* x *and its associated neural network parameters* θ*, whereas* $\mathcal{L}_{skewness}$ *is contingent upon the local input dataset* x' *and corresponding neural network parameters* θ*. Similarly, the latent space, akin to* $\mathcal{L}_{skewness}$*, is a function of the local input dataset* x' *and its respective NN parameters* θ*. As the local input dataset* x' *is invariably a subset of the global dataset* x*, these two sets were not differentiated in our notation of the total loss function. Consequently, we have revised the notation accordingly in Equations 2 (Page 4, Line 236-237), 3 (Page 4, Line 252-253), 4 (Page 4, Line 263-264), and 5 (Page 10, Line 644-645) within the main text.*

6. "A larger harmonic constant was applied when steering the reaction in the direction [...] to account for the higher barrier" how do you choose the harmonic constant? Is it difficult?

#Response#:

The reviewer raised an important point regarding the simulation setup. While this aspect is indeed relevant, it does not pose significant difficulties in practice. We realized that the original flowchart in Figure S1 did not accurately reflect the actual protocol used.

Rather than sequentially increasing the κ value after each exploration campaign, our approach involves launching a small swarm of parallel explorations, each with a different κ value, increasing within a specific range. This strategy ensures more efficient sampling of the reaction network while keeping computational cost and time under control.

We have now updated Figure S1 (Page 4, see SI) to reflect this parallel exploration scheme, which better represents the methodology we typically employ.

7. It's not clear how Loxodynamics finds multiple paths. Trial and error? When should I stop? Does this mean that the first path I find could be a higher FES pathway, and not the most common transition state?

#Response#:
We agree with the reviewer that, in its current form, Loxodynamics identifies multiple reaction pathways through a process that partially relies on trial and error. At its core, this process is strongly influenced by the choice of the κ parameter, which defines the strength of the biasing wall and thus guides the system toward the lowest-lying free energy pathways on the surface. Nonetheless, the procedure is highly automated and, as described in the reponse to comment 6 "this strategy ensures more efficient sampling of the reaction network while keeping computational cost and time under control".

To illustrate this point, we have included a series of computational experiments using the toy model PES presented in the main text, which feature two accessible reaction channels (Figure S6, Page 11, see SI). These examples help clarify how varying κ values can affect the exploration outcome.

Looking forward, we are actively working on integrating a reinforcement learning (RL) framework into Loxodynamics. In this approach, an agent will iteratively learn to discriminate and select the most promising directions among a swarm of independent trajectories. The objective is to increase the success rate of discovering low-barrier transition paths while maintaining diversity by discouraging convergence onto the same pathway.

Although still in its early stages, this RL-driven extension represents a key direction for the future development of Loxodynamics and will be the subject of forthcoming publications.

Remarks on code availability:

The code is clearly written and reasonably documented. It contains all the necessary tools to reproduce the results of the paper or use the proposed method on different systems, including a Jupyter notebook tutorial.

Reviewer #2

Remarks to the Author:

The manuscript by Zhang and Piccini presents an original and promising approach to reaction discovery based on the skewness of probability distribution using an autoencoder framework. The idea behind the method—Loxodynamics—is clearly explained and introduces a creative use of probability distribution

asymmetry to guide sampling. However, several methodological and validation aspects require clarification and improvement before the work can be recommended for publication.

Major Comments

≫ **Related Work and Literature.**

The paper lacks a thorough discussion of prior literature on reaction discovery using machine learning collective variables, especially autoencoder-based methods. Relevant prior work includes: We appreciate the reviewer's careful attention in highlighting fundamental literature that should be cited. Below, we outline the additions we have incorporated into the manuscript to address these suggestions.

- Chen & Ferguson https://doi.org/10.1063/1.5023804 demonstrated the use of autoencoders for interleaved CV discovery and enhanced sampling, iteratively expanding the sampled space.

#Response#:
We added the suggested reference to Chen and Ferguson J. Chem. Phys. one we already cited on the same topic (Page 3, Line 208-219, see main text).

- Ketkaew & Luber https://doi.org/10.1021/acs.jcim.2c00883 applied autoencoders to generate blind CVs for chemical reactions starting only from reactants.

#Response#:
We added the suggested reference (Page 3, Line 219-231, see main text).

- Raucci *et al.* https://doi.org/10.1021/acs.jpclett.1c03993 proposed a method for blind reaction discovery using general-purpose CVs with OPES explore.

#Response#:
We added the suggested reference (Page 2, Line 90-91, see main text).

Nonetheless, we believe that a further clarification on the distinction between Loxodynamics and other schemes is warranted. The aforementioned approaches employ CVs together with enhanced sampling methods such as umbrella sampling, metadynamics, or OPES to accelerate sampling. These methods, however, do not encode a specific directionality. In contrast, Loxodynamics introduces a bias explicitly designed to drive the system toward a new state along the CV. Metadynamics and OPES, in particular, are "flooding" approaches, in which the bias gradually fills the local free-energy basin to enhance fluctuations within it. The philosophy of Loxodynamics is fundamentally different: rather than amplifying fluctuations, it reshapes them with the specific purpose of propelling the system uphill.

≫ **Method Validation**
- Comparison with standard autoencoders: what advantage does the "Skewencoder" offer over a standard

autoencoder with the same architecture? A comparative experiment (e.g., on the toy model) would help clarify its contribution.

#Response#:

We thank the reviewer for raising this important point. Indeed, in our initial attempts, we employed a standard Autoencoder to achieve our goal. This involved projecting the configurations onto the latent space and evaluating the skewness of the resulting distribution without explicitly optimizing for it. In principle, such an approach could work.

Accordingly, we performed comparative experiments involving both the current Skewencoder configuration and a pure Autoencoder format (Section S2.A.1, Page 10-11, see SI), alongside evaluations of various formulations of the skewness loss term (refer to Reviewer #1's Point 4).

However, we consistently observed that, especially in the early stages of exploration when data are scarce, the skewness signal tends to be weak and significantly affected by noise. As a consequence, the direction selected by the biasing walls—based on this unreliable signal—amounts to an effectively random search (see Figure S5, Page 10, see SI). The bias then fails to consistently steer the system along meaningful reaction paths, greatly reducing the method's effectiveness.

These limitations motivated the design of the Skewencoder architecture, where the latent space is explicitly optimized to enhance the skewness signal, thereby providing a more robust and directional driving force for exploration.

- Minimum Free Energy Pathways: How can it be guaranteed that the discovered paths correspond to minimum free energy paths? Except for the first toy model, the others only have one pathway. The last example has two pathways, and the fact that it can also find a higher pathway seems to contradict this aspect. Also, it seems that this will depend on the hyperparameters (e.g., κ, ϵ)? Is prior knowledge of energy barriers required to set them?

#Response#:

We appreciate the reviewer's comment and we gladly introduced further information in the submitted revised version.

Firstly, we must clarify that κ is not a hyperparameter of the Skewencoder but rather a variable for the exploration process. As indicated in the revised Figure S1 in SI (Page 4, see SI) (also referenced in Reviewer #1's Point 6), no prior knowledge of κ is necessary. In the newly updated version of Section IV.C (Page 11, Line 712-731, see main text), we specify that instead of sequentially increasing κ after each exploration campaign, our method deploys a small swarm of parallel explorations, each with distinct κ values. This approach facilitates more efficient sampling of the reaction network while maintaining control over computational cost and time. Thus, κ is initialized as a parameter for agnostic exploration of Loxodynamics to thoroughly and efficiently search the chemical reaction space.

Concerning the influence of κ, we invite the reviewer to refer to Point 7 from Reviewer #1, where we detailed our evaluation of exploration performance by initializing various κ values within the Loxodynamics simulation applied to the double-well model PES illustrated in the main text. Briefly, a higher κ tends to identify reaction pathways with elevated reaction barriers, whereas a suitably lower κ is inclined to discover pathways with reduced reaction barriers, closely approximating the anticipated MEP. Thus, selecting κ can significantly affect the exploration of distinct reaction

pathways.

In the original manuscript, we had already discussed the role of both σ and ϵ in determining the position of the harmonic wall within the latent space (Page 10, Line 699-705, see main text). In the revised version, we have expanded the methodology section and included a new figure (Figure 7, Page 10, see main text). This figure illustrates that ϵ acts as a tunable offset, effectively controlling the degree of "brute force" applied during Loxodynamics simulations to explore new regions of the landscape.

Importantly, by adjusting the newly-added wall based on projected sample variance, this wall can explore unknown chemical spaces regardless of ϵ's value. Provided that ϵ is greater than 0, it will not constrain the range of extended chemical space exploration. Therefore, Loxodynamics' performance is not particularly sensitive to ϵ's selection, especially when it is small and positive. Since all sampling data are projected into the latent space and normalized—i.e., explored (known) space latent CVs are normalized to [-1, 1]—selecting ϵ at 0.1 represents a balanced choice that enhances simulation efficiency while preserving continuous and smooth exploration processes. The chosen ϵ remains consistent across all test cases, ensuring its neutrality concerning our method. In other words, controlling insensitive variables enables us to concentrate on analyzing κ's effect on Loxodynamics simulation performance.

Additionally, to ascertain whether the MD trajectory generated from Loxodynamics simulation approximates the MEP of the identified reaction pathway, we utilized the Climbing-Image Nudged Elastic Band (CINEB) method to optimize our MD trajectory snippet depicting reactions in both concerted and stepwise mechanisms (Section S2.C.3, Page 18, see SI). In summary, the geometries of the MD trajectory from Loxodynamics simulation closely resemble optimized configurational geometries on NEB paths, suggesting that the MD trajectory is already near the MEP of discovered reactions and indicating that Loxodynamics can guide systems approximately towards calculated MEP directions for certain reactions. Further simulation details are available in Section S2.C.3 (Page 18, see SI).

- The discussion section claims a wide range of potential applications, but the validation is limited to a few relatively simple cases (except for the last one). A more realistic assessment of current limitations, along with suggestions for how to transition from exploration to detailed mechanistic modeling, would strengthen the manuscript. For instance: after discovery, can the CVs produced by Skewencoder be used in conventional biased simulations to compute free energy surfaces?

#Response#:

We appreciate the reviewer for this insightful comment. In the newly added Section S2.B.3 (Page 15-16, see SI), we present results from well-tempered Metadynamics simulations employing the learned Skewencoder CV in both S_N2 and Diels-Alder reactions. These Skewencoder CVs are derived from iterations of Loxodynamics simulations where state transitions occurred. By properly configuring Metadynamics parameters, Figure S11 (Page 16, see SI) demonstrates that the Skewencoder CV can facilitate transitions between two states in Metadynamics simulations of both reactions. The Skewencoder CV also effectively differentiates reactant and product states, corresponding to the state information obtainable from heuristic CVs. Figure S12 (Page 16, see SI) illustrates that samples from well-tempered Metadynamics using Skewencoder CVs can reconstruct the FES for specific reactions. The reconstructed FES aligns with those generated by Metadynamics simulations employing heuristic CVs. Consequently, Skewencoder CVs can be further applied in conventional biased simulations, provided they encapsulate configurational information within both reactant and product spaces.

Moreover, we have included a new test case involving a significantly more complex chemical reaction network, which has been explored in an agnostic manner using Loxodynamics. Please check our answer to comment 4 to Reviewer #4.

≫ **Methodological Details**

- It is unclear whether κ is chosen by the user or adapted during the iterations (as suggested. In the SN2 example, it is said: "Due to the asymmetric energy barriers… a larger harmonic constant was applied … to account for the higher barrier." Does setting this parameter require prior barrier knowledge?

#Response#:

We thank the reviewer for raising this point. In practice, we initiate a small swarm of independent trajectories (typically 5 to 10 depending on the complexity of the problem), each characterized by a different value of the biasing parameter κ. This strategy allows us to explore multiple directions efficiently and identify suitable paths for further exploration.

We acknowledge that this detail was not clearly explained in our original submission. To address this, we have updated both the main manuscript (Section IV.C, Page 11, see main text) and the Supporting Information (Figure S1, Pages 4, see SI), where we now explicitly describe this initial swarm approach. Additionally, we revised the flowchart in Figure S1 to reflect this procedure, clearly indicating the use of an initial set of test trajectories that guide the subsequent path exploration.

In general, no precise prior knowledge of the barrier is required.

- The rationale behind the inclusion of σ in the wall position (i.e., μ + σ + ϵ) is not explained in the main text. Figure 1 should also reflect this logic explicitly. Also, the role of ϵ seems important for determining the degree of exploration, but is not discussed. A discussion of how to choose these parameters and what effect they have is needed.

#Response#:

We appreciate the reviewer for highlighting this issue. In the newly revised manuscript, we provide a detailed explanation in the main text regarding the impact of utilizing projected sample variance σ and a custom offset ϵ, as illustrated in Figure 7 (Page 10, see main text), rather than merely portraying them in Figure 1 as before.

In the current revised version, Figure 1 (Page 3, see main text) primarily illustrates that the relationship between sample skewness and local anharmonicity arises from biased simulations on an artificial 1-D PES. During each iteration of the 1-D simulation, a bias harmonic wall is introduced and consistently shifted with a constant offset Δx from the left to the right along the x-axis. We observed that as the system transitioned from one minimum to another, the sample distribution became increasingly asymmetric, reflecting its gradual evolution towards state transition. In the current manuscript version, we have clarified that this observation from the 1-D simulation served as inspiration for going from a simple ideal case to more realistic scenarios. Consequently, to address challenges in extending this knowledge to multidimensional scenarios, we developed the Skewencoder NN architecture and Loxodynamics workflow as innovative tools for chemical space exploration.

The impact of the custom offset ϵ has been addressed in the response to the Method Validation section of Major Comments from Reviewer #2. Similarly, we employ projected sample variance to determine the harmonic wall's position.

Calculating projected sample skewness necessitates first computing sample variance; thus, incorporating this term into our algorithm incurs no additional computational cost. Furthermore, sample variance reflects the distribution width within latent space. By using this term as an offset for the harmonic wall, previously explored regions in latent space are bypassed in new simulation iterations, reducing overlap and enhancing exploration efficiency. In practice, one may replace projected sample variance σ or even σ combined with ϵ with a constant offset (as illustrated in Figure 1, page 3, see main text). This substitution yields varying effects across different PES. By introducing σ and a moderate ϵ, we devised a mechanism that adaptively adjusts Loxodynamics' exploration "speed" based on local datasets without incurring extra computational expense.

- The claim that the reaction discovery is "unbiased" is misleading. Since the skewness of the biased distribution depends on the applied parameters of the harmonic bias and no reweighting is performed, the method introduces a directional sampling artifact.

#Response#:

We appreciate the reviewer for highlighting this issue. We would like to clarify that we do not claim the reaction discovery process to be ***unbiased****, but rather* ***agnostic****. Although it is true that no reweighting is performed, such an action would primarily reproduce the equilibrium distribution of the initial state, providing limited insight into transition pathways. Therefore, a certain degree of directionality must be introduced. Importantly, this directional bias is not imposed a priori by the user but emerges during the early stages of the simulation when the system experiences minimal or no bias. As exploration progresses, the method selects specific directions for sampling reaction pathways. This selection is not manually adjusted but driven by an iterative sampling protocol that dynamically adapts to the underlying free energy landscape. For clarity, we have removed "unbiased" from the sentence: By following pathways with the lowest energy barriers, this approach enables efficient and unbiased exploration of likely reaction channels (Page 2, Line 101, see main text).*

- A complex procedure to decide whether to warm-start the autoencoder or not is put forward, but not justified. A comparison with an exploration training only from scratch on the toy model would clarify its impact. Also, this procedure requires a criterion to determine whether a new metastable state has been reached, which is not described for the examples reported. This contradicts the claim of minimal user intervention and should be detailed more clearly.

#Response#:

We appreciate the reviewer for raising this insightful issue.

In the revised version of our manuscript, we have included Section S1.C (Page 5, see SI) to elucidate the criterion for distinguishing between metastable states and non-metastable states. Generally, this criterion can be adapted to any connectivity matrix-based method capable of determining atomic connectivity through customized switching functions. The present implementation is an in-house simplified version designed for efficiency.

Additionally, a comparative experiment was conducted between the Loxodynamics simulation utilizing the warm start training strategy and exploration training initiated from scratch on the toy model PES. The results are presented in Section S2.A.3 (Page 11, see SI). In summary, the warm start training strategy enhances exploration efficiency by lowering training and computational costs, thereby fulfilling the intended objective of this strategy within Loxodynamics

simulation.

- Section A of the methods could benefit from a few revisions. I think that the workings of neural network layers are now part of common knowledge and so can be moved to the SI.

#Response#:
We thank the reviewer for this valuable suggestion. This content has been relocated to Section S1.B.2 (Page 3, Line 43-51, in SI), integrated with the discussion on the NN structure of Skewencoder, in the newly revised version of the manuscript.

On the other hand, it is not clear if the details about the $\mathcal{L}_{\mathtt{skew}}$ are just conceptual considerations (as it seems to me looking at the code) or technical details that need to be implemented to improve the performance of the method. Whichever the case, I think this should be made clear to the reader.

#Response#:
We appreciate the reviewer for identifying this issue. The current format of $\mathcal{L}_{skew}$ encompasses both conceptual considerations and technical implementation, enhancing the method's performance. Refer to Point 4 from Reviewer #1 and the discussion in the newly added Section S2.A.1 (Page 8-9, in SI) regarding this matter. In summary, the current implementation of $\mathcal{L}_{skew}$ has been shown to improve exploration efficiency when compared with other forms of $\mathcal{L}_{skew}$. Additionally, Skewencoder utilizing $\mathcal{L}_{skew}$ in a purely Gaussian-like form successfully identified alternative states and exhibited only slightly reduced performance during Loxodynamics simulation on the toy PES, suggesting that $\mathcal{L}_{skew}$ can be applied flexibly and is therefore also a conceptual consideration.

- In general, it would be much better to include in the methods section a summary of the method, including the loss functions (which are only defined in the introduction) and the workflow of the iterative procedure (which is again discussed in the introduction and summarized only in the SI).

#Response#:
We appreciate the reviewer for this helpful suggestion. We have reorganized the methodology section in the main text accordingly. To eliminate redundancy, we continue to focus on discussing the design of the skewness loss-related term in Section IV.A (Page 9-10, see main text), but have removed the introduction to a general NN structure. Subsequently, we partially revised Section VI.B (Page 10, see main text). We included Figure 7 (Page 10, see main text), which more clearly illustrates how the harmonic wall is applied based on sample mean, variance, and skewness values. Additionally, we employed a simplified equation form indicating how the harmonic wall is applied in latent space, as suggested by Reviewer #1 in Point 3. As a conclusion to the methodological section, we present a simplified version of the entire algorithm in Section IV.C (Page 11, see main text), corresponding to detailed discussions in Section S1 of the SI.

Comments specific to the applications presented.

$\gg$ **Toy Model**
- The evolution of the skewness and the mean of the CV across iterations is not shown. Including this would help illustrate how the method progresses and adapts the biasing direction based on the measured skewness?

#Response#:
We appreciate the reviewer for highlighting this issue. We wish to clarify that the latent space undergoes changes at each iteration, rendering comparisons of distribution moments across different variables impractical.

Concerning CV evolution as requested by Reviewer #1, please refer to Point 1 from Reviewer #1 where we illustrate the evolution of CVs in toy models.

- How does the exploration change as a function of the chosen hyperparameters, κ and ϵ? Showing supplementary examples would help clarify how to set them.

#Response#:
We appreciate the reviewer for this comment. We have previously addressed the impact of ϵ in earlier questions (e.g., Method Validation section of Major Comments from Reviewer #2). Regarding the effect of κ, as discussed in previous feedback, briefly, a higher κ tends to identify reaction pathways with elevated reaction barriers, whereas a suitably lower κ is inclined to discover pathways with reduced reaction barriers, closely approximating the anticipated MEP. Thus, selecting κ can significantly influence the exploration of distinct reaction pathways. We have added Section S2.A.2 (Page 9-11, see SI) to assess the effect of κ using a toy model in Loxodynamics.

Concerning Skewencoder hyperparameters, we assert that with a typical deep autoencoder neural network structure, Skewencoder can effectively learn features from latent space and drive Loxodynamics. We have already compared Skewencoder performance using both deeper and shallower neural network structures in Section S2.B.2 (Page 13-15, see SI). Both configurations of the neural network successfully drive the system to identify state transitions. Therefore, given the robust fitting capability of neural networks, a general Autoencoder structure can be employed within Skewencoder without extensive hyperparameter optimization efforts. This highlights our method's flexibility and scalability.

- Although the potential includes two distinct paths, only one is sampled. Can the method recover both, depending on the hyperparameters used? This would also help clarify how this is obtained in the chabazite case.

#Response#:
We appreciate the reviewer for highlighting this issue. We kindly request the reviewer to refer to Point 7 from Reviewer #1 and Section S2.A.2 (Page 9-11, see SI).

$\gg$ SN2 Reaction
- The input CV space essentially consists of d_{Cl} and d_F, which already capture the transition without needing any combination of them. In my opinion, this system should be moved to the SI, as it does not significantly validate the method.

#Response#:
We appreciate the reviewer for this suggestion. However, we would like to emphasize that although the S_N2 reaction is a relatively simple example, it serves as a realistic test case. Given that the input dimensions of the previous model

PES are all 2D, and the S_N2 reaction involves two Skewencoder input descriptors, it can be considered an appropriate intermediary between more complex real-world examples and simpler model PES, which are easier to test. Consequently, we have decided to retain this section in the main text, as it continues to validate the method.

≫ **Chabazite**

- How is the second pathway found? In the SI it is said that different elastic constants are used to reflect barrier differences. Since barrier information is usually unknown, how would one proceed in general?

#Response#:

This is certainly a fundamental point. As discussed in our response to Point 6 from Reviewer #1 and illustrated in Figure S1 (Page 4, see SI), our method initiates a small swarm of parallel explorations, each with a distinct κ value that increases within a specified range. This approach allows Loxodynamics simulation to identify corresponding reaction pathways with different κ values without requiring barrier information. In the SI, we mention that various constants are employed to reflect barrier differences, indicating that different κ values configuring harmonic walls correspond to barrier height differences calculated using DFT methods (as reported in https://pubs.acs.org/doi/10.1021/jp513024z). We do not imply that κ directly corresponds to energy values; rather, smaller κ values are likely to uncover reaction pathways with lower barriers, whereas larger κ values tend to reveal pathways with higher barriers. This observation aligns with our findings in the double pathway model PES presented in Section S2.A.2 (Page 9-11, see SI).

- In the stepwise mechanism, how is backward motion prevented? Are additional walls required? Would this become limiting for more complex catalytic cycles involving many intermediate steps?

#Response#:

We appreciate this comment very much as we realize this point was not clear in the original submission. Thus, we updated accordingly Section S2.C.2(Page 18, see SI), discussing how we restart the Loxodynamics simulation to explore a reaction pathway that includes multiple reaction wells, e.g. the stepwise mechanism in the ethanol dehydration case. In brief, Figure S13 (Page 17, see SI) shows that once we reached the intermediate state, we restart the Loxodynamcis, keep all the parameters related to the Skewencoder, but we retain the old CV along which we apply an additional restraining wall that guide the system to identify the current intermediate state. Then in the newly restarted Loxodynamics simulation, an "unbiased" simulation will be conducted, with the restraining wall kept. Then new Loxodynamics bias will be added in the following iteration, with the restraining wall still existing. To compensate the additional potential bias caused by the restraining wall, κ value must be increased to explore the second potential well, though the barrier height of the first reaction and the second reaction in the multiple-well reaction path are similar. We reported this observation in Section S2.C.2 (Page 18, see SI) and S2.D (Page 18-20, see SI), corresponding to different dehydration reactions on the acidic Chabasite. Based on our observation, we assume that this might be a limitation to a reaction pathway with more wells. However, in the test cases we have conducted Loxodynamics simulations on so far we do not observe such a limitation.

≫ **Minor Comments**

- Many figure captions are not self-contained. For example, the meaning of color schemes and symbols (e.g., Figure 1 symbols, Figure 3 colorbar, color of scatter plots, Figure 4 color of scatter plots, Figure S1 symbols, Figure S2 colorbars, Figure S4 color of scatter plots, colorbars) is only defined in the main text

and should be reported in the captions.

#Response#:

We appreciate the reviewer for identifying this issue. Below, we detail the content added to the captions of the relevant figures to address this concern. Generally, we emphasize that the precise values on the colorbar (Figure 3 and Figure S2) for both the model PES and projected latent space (Figure S4) are not critical. The closed format has already been provided for the PES. As Loxodynamics simulations progress, the projected CV space evolves continuously, with exact values in projected latent space CV being updated iteratively. Consequently, relative relationships are more significant than exact values for clearly illustrating CV evolution.

In the caption of Figure 1 (Page 3, see main text), we added:

- The initial wall is located at $x = \mu$, corresponding to the mean of the unbiased sample distribution (indicated by red dots in the upper panel). In each subsequent iteration, the harmonic wall is translated by a fixed offset Δx relative to its position in the previous iteration. The spring constant κ for all walls (depicted as blue, green, and yellow half-parabolic curves in the lower panel) remains constant throughout all iterations.

In the caption of Figure 3 (Page 5, see main text), we added:

- The darker region signifies an area with lower potential values, whereas the lighter region denotes an area with higher potential values.

- The central and lower panels display the evolution of the x (in dark green) and y (in light green) coordinates, respectively, over time (in Lennard-Jones units). The blue points represent samples from the incremental dataset, while red points correspond to samples from the current iteration.

In the caption of Figure 4 (Page 6, see main text), we added:

- In the 2D maps, as in Figure 3, blue points denote samples from the global dataset, whereas red points indicate samples from the local dataset corresponding to the current iteration. Relevant structures extracted from the trajectory blocks of each successive iteration are shown in the upper subpanels.

In the caption of Figure S1(Page 4, see SI), we added:

- In this flowchart, $\{\kappa_i\}$ represents a sequence of wall spring constants arranged in an arithmetic progression with a common difference of $\delta\kappa$. The parameter n_{iter} denotes the maximum number of iterations permitted for the Loxodynamics simulation. The quantity γ_{ls} refers to the projected sample skewness in latent space. The variable $\mathbf{L}_{state}$ indicates the list of identified metastable states, while S_p serves as the identifier for the current state.

In the caption of Figure S3 (Page 7, see SI), we added:

- The darker region signifies an area with lower potential values, whereas the lighter region denotes an area with higher potential values.

In the caption of Figure S4 (Page 8, see SI), we added:

- In each left subcolumn, The blue points represent samples from the incremental dataset, while red points correspond to samples from the current iteration.

- Lighter regions correspond to higher values in the latent CV space, whereas darker regions represent lower values.

- The provided code is well-structured and detailed, which ensures reproducibility. However, from the SI and the code, it appears to rely extensively on the mlcolvar library, which is not cited in the main text.

#Response#:
We appreciate the reviewer for this suggestion. The relevant citation has been incorporated in the code and data availability section (Page 12, see main text).

- When quantities are discussed in the main text, refer to the specific equation number to help guide the reader (e.g., $\mathcal{L}_{\mathtt{skew}}$ in methods sec. A).

#Response#:
We appreciate the reviewer for this suggestion. We have added the following in the main text (Page 9-11, line 631-633): This section is devoted to the discussion of the formulation and implementation of the skewness-related loss term, $\mathcal{L}_{skew}$ (defined in Equation 4), to address this issue.

Remarks on code availability:

The provided code is well-structured and detailed, which ensures reproducibility. However, from the SI and the code, it appears to rely extensively on the mlcolvar library, which is not cited in the main text.

Reviewer #3

Remarks to the Author:

I co-reviewed this manuscript with one of the reviewers who provided the listed reports. This is part of the Nature Communications initiative to facilitate training in peer review and to provide appropriate recognition for Early Career Researchers who co-review manuscripts.

Reviewer #4

Remarks to the Author:

In this manuscript, Zhang and Piccini propose a novel approach to explore reactive paths using an encoder that includes the skewness of the probability distribution observed in a (set of) collective variables into its loss function. I find the manuscript to be clearly written and the method is well-explained. The authors, moreover, support their method with a handful of examples ranging from toy potentials to heterogeneous catalysis. While I would most certainly recommend publication in a more specialized journal, I believe that, in their present form, the results might not be sufficiently groundbreaking or of broad enough applicability

given the multidisciplinary audience of Nature Communications. My main points of concern, in this regard, are the following:

1. The authors state "Moreover, upon algorithm termination, the accumulated samples provide sufficient data to construct a latent space that effectively approximates a meaningful CV". While it looks like that is the case with the simple toy systems, it would be interesting to see if the conclusion holds when going to the more realistic reactions (e.g., plotting the isolines on the d_6 vs d_4 graph in Figure 4). I would find a reaction finder that also provides an (ok-ish) CV to sample the reaction with enhanced sampling quite interesting.

#Response#:
We appreciate the reviewer for highlighting this issue. As pointed out also in our response to Reviewer #2, in the newly added Section S2.B.3 (Page 15-16, see SI), we present results from well-tempered Metadynamics simulations employing the learned Skewencoder CV in both S_N2 and Diels-Alder reactions. These Skewencoder CVs are derived from iterations of Loxodynamics simulations where state transitions occurred. By properly configuring Metadynamics parameters, Figure S11 (Page 16, see SI) demonstrates that the Skewencoder CV can facilitate transitions between two states in Metadynamics simulations of both reactions. The Skewencoder CV also effectively differentiates reactant and product states, corresponding to the state information obtainable from heuristic CVs. Figure S12 (Page 16, see SI) illustrates that samples from well-tempered Metadynamics using Skewencoder CVs can reconstruct the FES for specific reactions. The reconstructed FES aligns with those generated by Metadynamics simulations employing heuristic CVs. Consequently, Skewencoder CVs can be further applied in conventional biased simulations, provided they encapsulate configurational information within both reactant and product spaces.

2. Concerning the last case study, I find it presently unclear what the difference was between the simulations that led to the exploration of the concerted vs stepwise mechanisms (aside from the wall kappa, as reported in the SI). Was sampling the two paths in the two simulations a coincidence?

#Response#:
We appreciate the reviewer for this comment which has also been raised by other Reviewers.

The sampling of two paths is not coincidental but rather contingent upon κ as a parameter initialized within a swarm. As discussed in our response to Point 6 from Reviewer #1 and illustrated in Figure S1 (Page 4, see SI), our method initiates a small swarm of parallel explorations, each with a distinct κ value that increases within a specified range. This approach allows Loxodynamics simulation to identify corresponding reaction pathways with different κ values without requiring barrier information. In the SI, we mention that various constants are employed to reflect barrier differences, indicating that different κ values configuring harmonic walls correspond to barrier height differences calculated using DFT methods (as reported in https://pubs.acs.org/doi/10.1021/jp513024z). We do not imply that κ directly corresponds to energy values; rather, smaller κ values are likely to uncover reaction pathways with lower barriers, whereas larger κ values tend to reveal pathways with higher barriers. This observation aligns with our findings in the double pathway model PES presented in Section S2.A.2 (Page 9-11, see SI).

How different are they in energy according to the semiempirical method used by the authors?

#Response#:

We appreciate the reviewer for highlighting this insightful issue. To evaluate how closely the trajectory obtained from Loxodynamics simulation approximates the MEP of a specific reaction, particularly ethanol dehydration at acidic Chabazite, we employed the Climbing-Image Nudged Elastic Band (CI-NEB) method to optimize our MD trajectory snippet depicting reactions in both concerted and stepwise mechanisms. We introduced a refined CI-NEB workflow in Section S2.C.3 (Page 18, see SI), newly added. This workflow begins with conducting CI-NEB calculations using a relatively long MD trajectory snippet as input, followed by selecting representative images from the initial calculation results to serve as initial images for a second CI-NEB calculation. This approach allows the refined CI-NEB workflow to utilize information from the Loxodynamics simulation trajectory while reducing computational costs. The NEB energies calculated through this method are comparable to DFT calculation results (as reported in https://pubs.acs.org/doi/10.1021/jp513024z). Furthermore, the geometries of the MD trajectory closely resemble optimized configurational geometries on NEB paths, indicating that the MD trajectory used as input is already near the MEP of discovered reactions and suggesting that Loxodynamics can guide systems approximately towards calculated MEP directions for certain reactions. Further simulation details can be found in Section S2.C.3 (Page 18, see SI)

Also, in the stepwise mechanism, how likely it is to actually resample the backward reaction from the surface species to ethene instead of going forward?

#Response#:

We appreciate the reviewer for highlighting this issue.

As pointed out also in our response to Reviewer #2, we realize this point was not clear in the original submission. Thus, we updated accordingly Section S2.C.2(Page 18, see SI), discussing how we restart the Loxodynamics simulation to explore a reaction pathway that includes multiple reaction wells, e.g. the stepwise mechanism in the ethanol dehydration case. In brief, Figure S13 (Page 17, see SI) shows that once we reached the intermediate state, we restart the Loxodynamcis, keep all the parameters related to the Skewencoder, but we retain the old CV along which we apply an additional restraining wall that guide the system to identify the current intermediate state. Then in the newly restarted Loxodynamics simulation, an "unbiased" simulation will be conducted, with the restraining wall kept. Then new Loxodynamics bias will be added in the following iteration, with the restraining wall still existing. To compensate the additional potential bias caused by the restraining wall, κ value must be increased to explore the second potential well, though the barrier height of the first reaction and the second reaction in the multiple-well reaction path are similar. We reported this observation in Section S2.C.2 (Page 18, see SI) and S2.D (Page 18-20, see SI), corresponding to different dehydration reactions on the acidic Chabasite. Based on our observation, we assume that this might be a limitation to a reaction pathway with more wells. However, in the test cases we have conducted Loxodynamics simulations on so far we do not observe such a limitation.

3. I find that some comparison with literature is missing. For example, how does the method proposed by the authors compare with OPES-based exploration methods (e.g., https://pubs.acs.org/doi/10.1021/acs.jpclett.1c03993) in terms of cost and ability of finding new reaction products?

#Response#:

We appreciate the reviewer for highlighting this point. We have now included a sentence in the main text (Page 2, Line

89-90, see main text) explicitly comparing Loxodynamics to OPES_E: "In general, designing a fully agnostic CV for reactivity exploration with high sampling efficiency is impractical without prior knowledge of the expected products."

While OPES_E can identify new metastable states, it still depends on an initially defined CV, necessary for applying bias. Additionally, OPES_E—similar to other flooding techniques like metadynamics or variationally enhanced sampling—operates by gradually filling free energy wells along the predefined CV, enabling transitions to other metastable states.

Conversely, our method resembles directional pulling approaches such as steered MD or ratchet-and-pawl MD, where bias is applied in a specific direction. The primary advantage of Loxodynamics is that both the CV and pulling direction are not predetermined by the user but are dynamically and automatically updated during exploration. This facilitates more agnostic and adaptive sampling of reaction pathways.

4. The reported case studies, while comprehensive, are all pretty simplistic. To remain in zeolite catalysis, I guess that using a larger adsorbed molecule (with, e.g., double bonds) could open up quite a lot more reactive paths (cyclization, cracking, and so on). Going back to comment 2, it would be interesting to see which paths is the method able to discover.

#Response#:

We appreciate the reviewer for the insightful suggestion. Therefore, we embraced the suggestion and started a much more complex case, with a larger reaction network. More precisely, in the newly revised manuscript, we introduce a more complex example in zeolite catalysis, namely 1-Butanol Dehydration in acidic chabazite.

The 1-Butanol case shares similar characteristics with ethanol dehydration, consisting of two reaction mechanisms—concerted and stepwise pathways. A key distinction between 1-butanol and ethanol is the diverse array of potential reaction products, influenced by double bond positioning along the alkyl chain and associated isomerism. Literature indicates that 1-butanol dehydration can yield various products and may be accompanied by side reactions like oligomerization, cracking, and aromatization under high substrate loading conditions. In the revised manuscript, we focus on single-site adsorption specifically examining dehydration pathways. Additional details are provided in Section II.D (Page 8-9, see main text).

Figure 6 (Page 8, see main text) depicts reaction pathways identified via Loxodynamics simulations. Along the direct dehydration route, the catalyst produces three canonical products of 1-butanol dehydration: 1-butene and cis- and trans-isomers of 2-butene. As expected, 1-butene formation may also proceed through stepwise mechanisms akin to those in ethanol dehydration pathways. Further detailed analysis is available in Section II.D (Page 8-9, see main text). The κ values reported in S2.D (Page 18-20, see SI) correspond to magnitude relationships among different products, particularly conformers.

This case study demonstrates our method's capability to efficiently explore a larger catalytic reaction network while identifying both likely products and pathways involving intermediates with comparable occurrence likelihoods. We believe this case study sufficiently illustrates Loxodynamics' potential application for exploring an extended realm of reaction spaces.

If the authors can reply to these comments, the manuscript might become suitable for Nature Communi-

cations.

I additionally have the following minor comment:
1. I am not fully sure I understand all the passages behind warm-start training explained in the SI. Am I correct in stating that warm start is used unless the system falls into a new basin of attraction, in which case everything restarts from scratch? In the final part, the authors mention the case in which the system falls back to the original minimum. I do not immediately see how training from scratch would suddenly push the system towards another barrier.

#Response#:
We appreciate the reviewer for highlighting this issue. We would like to clarify that the warm start training strategy is employed unless: 1) the system enters a new basin of attraction, or 2) the system reverts to the original reactant space without identifying another state. In the latter scenario, the system should transition from "unstable" states back to its original state. Definitions regarding whether a state is (meta)stable or unstable are provided in Section IV.D.2 (Page 11, see main text) and the newly added Section S1.C (Page 5, see SI).

In the scenario mentioned by the reviewer in Section S2.B.2 (Page 13-15, see SI), the system was abruptly reverted to its original state because the κ *value was sufficiently large to drive the system back and forth, in absence of additional restraining walls to prevent this behavior. This observation aligns with the shape of the FES obtained from Metadynamics simulations using either heuristic CV or Skewencoder CV for Diels-Alder reaction (Figure S12, Page 16, see SI). Figure S9 (Page 14, see SI) illustrates that a large* κ *forcefully pushes the system from a wide but higher basin to a lower and narrower basin, compared with trajectories generated by Loxodynamics simulation featuring deeper Skewencoder structures (Figure 4, Page 6, see main text). Despite training in the new state starting entirely from scratch, Skewencoder can still capture information from already sampled configurational space. Moreover, Loxodynamics harmonic wall bias is adaptively positioned using sample mean and variance values (Figure 7, Page 10, see main text), preventing redundant exploration of known configurational space. Under these conditions, the system can swiftly exit shallow basins and return to original states. In summary, here we aim to demonstrate that Skewencoder's structure is flexible and capable of training representative latent space CVs, underscoring our workflow's conceptual robustness. We acknowledge potential misconceptions caused by this example; hence we analyze warm-start training strategy effects on model PES in Section S2.A.3 (Page 11-12, see SI).*

Response to the reviewer's comments to manuscript

NCOMMS-25-28852A

"Exploring Chemistry and Catalysis by Biasing Skewed Distributions via Deep Learning"

Zhikun Zhang, GiovanniMaria Piccini

December 22, 2025

We sincerely thank the reviewers for their positive assessment of our revised manuscript and their acknowledgement of the improvements made in the previous round. We appreciate the additional suggestions provided, which have been very helpful in further polishing the text and clarifying the final details. We have addressed these remaining points below and believe the manuscript is now fully ready for publication. All additional suggestions from the reviewers are marked in red, while our corresponding responses are in blue italics.

Reviewer #1

Remarks to the Author:

The authors addressed all my concerns and the new version of the paper is in excellent shape, especially the much improved SI. I believe it makes a significant contribution to the field and I recommend it for publication.

Reviewer #2

Remarks to the Author:

I appreciate the Authors' efforts in addressing the reviewers' questions, which have clarified several important aspects of the work. That said, I believe that some minor aspects could still be refined to further improve the manuscript.

In particular, in several cases the clarifications provided in the response letter have not been incorporated into the manuscript, or have been included only in the Supporting Information without being mentioned in the main text. This makes it difficult for readers to access or appreciate these details. Since some of these issues were raised independently by multiple reviewers, adding the corresponding explanations or cross-references in the manuscript would help make the work clearer and more robust.

#Response#:
We thank the reviewer for these general comments. We have incorporated the corresponding explanations from the original response letter into the revised manuscript. For a detailed discussion, please refer to our responses to the specific remarks below.

Furthermore, the inclusion of the more complex example (1-butanol dehydration) is a welcome improvement and represents a valuable extension of the study. However, this example lacks sufficient supporting data: In addition to the schematic of the "discovered" species, some quantitative evidence (e.g., time-series data or scatter plots) should be reported in the SI.

#Response#:
We thank the reviewer for this comment. We have included Figures S15 and S16 in the Supporting Information (SI), which illustrate the evolution of the system using representative bond distances, analogous to the approach in Figure 5. The corresponding trajectories are available in the dataset published on Zenodo.

Regarding Section D.2, it would also be useful to briefly mention alternative or complementary approaches for state identification, as suggested by other reviewers; for instance, the use of machine-learning-based classifiers that could offer a more general or automated strategy.

#Response#:
We thank the reviewer for this insightful suggestion. We have added the following content to Section D.2 of the manuscript to provide a complete elaboration of the state detection method.

While the implemented state detection method is straightforward and effective for the presented case studies, complex systems may necessitate more robust approaches. State-of-the-art machine learning classifiers, particularly those designed for transition state detection, offer a generalized and automated strategy. These methods could provide the fast and precise state determination required for the warm-start training of Loxodynamics. Alternatively, for simpler organic fragments or structures that can be chemically isolated from their environment, SMILES-based fingerprinting serves as a valuable tool, allowing changes in bond topology to be rapidly identified via string analysis.

Additional Points Related to the Original Remarks

Major Comments

≫ **Related Work and Literature.**

The paper lacks a thorough discussion of prior literature on reaction discovery using machine learning collective variables, especially autoencoder-based methods. Relevant prior work includes: We appreciate the reviewer's careful attention in highlighting fundamental literature that should be cited. Below, we

outline the additions we have incorporated into the manuscript to address these suggestions.

- Chen & Ferguson https://doi.org/10.1063/1.5023804 demonstrated the use of autoencoders for interleaved CV discovery and enhanced sampling, iteratively expanding the sampled space.

#Response#:
We added the suggested reference to Chen and Ferguson J. Chem. Phys. one we already cited on the same topic (Page 3, Line 208-219, see main text).

- Ketkaew & Luber https://doi.org/10.1021/acs.jcim.2c00883 applied autoencoders to generate blind CVs for chemical reactions starting only from reactants.

#Response#:
We added the suggested reference (Page 3, Line 219-231, see main text).

Reviewer: The way this reference is mentioned is not contextualized correctly; it seems to even imply that it already solved the problem addressed in the present work by the Authors, which contrasts with the clarification given below.

#Response#:
We appreciate the reviewer for highlighting this issue. We have reformulated the discussion in the manuscript to clarify that the primary advantage of Loxodynamics is its ability to actively drive exploration toward new regions of phase space via an explicit directional bias. Notably, our method achieves this without relying on conventional enhanced sampling paradigms or static CVs derived from extensive prior simulations.

"Ketkaew et al. introduced an autoencoder-based strategy aimed at constructing data-driven collective variables for enhanced sampling, trained exclusively on reactant-state configurations. Their approach exploits deep autoencoders to extract latent representations from unbiased trajectories and employs these representations within flooding-based enhanced sampling schemes to reconstruct free energy surfaces. While this methodology alleviates the need for explicit product-state information in the definition of collective variables, it remains inherently tied to static CVs and relies on conventional enhanced sampling paradigms, without introducing an explicit directional bias capable of actively driving exploration toward new regions of phase space."

- Raucci *et al.* https://doi.org/10.1021/acs.jpclett.1c03993 proposed a method for blind reaction discovery using general-purpose CVs with OPES explore.

#Response#:
We added the suggested reference (Page 2, Line 90-91, see main text).

Reviewer: This reference is cited but even less contextualized. Considering that it addresses the same issue the Authors discuss, it should be at least mentioned in the part about similar approaches; together with the clarification below.

#Response#:
We appreciate the reviewer for raising this point. Raucci et al. proposed an exploration strategy based on Social Permutation Invariant Coordinates (SPRINT) combined with the OPES method. However, their single-dimensional CV—the largest eigenvalue of the connectivity matrix—depends strictly on the initial geometry. Consequently, the search direction is effectively predetermined at the start. In contrast, Loxodynamics explores chemical space dynamically; by varying the wall constant κ, different product spaces can be accessed. Furthermore, due to its distinct sampling pattern, Loxodynamics demonstrates superior efficiency, driving the system to new states with five to ten times less simulation time.

Regarding the reaction coordinate definition, Raucci et al. employ Deep LDA to refine the CV describing the reactant and product only after a new state is detected. Conversely, Loxodynamics yields a CV during the exploration phase itself, which is often immediately suitable for subsequent enhanced sampling methods like Metadynamics.

We initially omitted a detailed discussion of this paper because it is not an autoencoder-based method. Instead, it falls under the category of methods utilizing predefined CVs with conventional enhanced sampling—approaches which, as we noted, often lack the flexibility to sample diverse directions. However, to provide a complete context, we have incorporated the following explanation into the main text and added the citation for the reader's reference.

Nonetheless, we believe that a further clarification on the distinction between Loxodynamics and other schemes is warranted. The aforementioned approaches employ CVs together with enhanced sampling methods such as umbrella sampling, metadynamics, or OPES to accelerate sampling. These methods, however, do not encode a specific directionality. In contrast, Loxodynamics introduces a bias explicitly designed to drive the system toward a new state along the CV. Metadynamics and OPES, in particular, are "flooding" approaches, in which the bias gradually fills the local free-energy basin to enhance fluctuations within it. The philosophy of Loxodynamics is fundamentally different: rather than amplifying fluctuations, it reshapes them with the specific purpose of propelling the system uphill.

Reviewer: This answer clarifies the difference with respect to the other methods. Adding it to the manuscript would be highly beneficial.

#Response#:
We thank the reviewer for this suggestion. As noted in the previous point, we have incorporated the relevant discussion into the main text.

≫ **Method Validation**
- Comparison with standard autoencoders: what advantage does the "Skewencoder" offer over a standard

autoencoder with the same architecture? A comparative experiment (e.g., on the toy model) would help clarify its contribution.

#Response#:
We thank the reviewer for raising this important point. Indeed, in our initial attempts, we employed a standard Autoencoder to achieve our goal. This involved projecting the configurations onto the latent space and evaluating the skewness of the resulting distribution without explicitly optimizing for it. In principle, such an approach could work.

Accordingly, we performed comparative experiments involving both the current Skewencoder configuration and a pure Autoencoder format (Section S2.A.1, Page 10-11, see SI), alongside evaluations of various formulations of the skewness loss term (refer to Reviewer #1's Point 4).

However, we consistently observed that, especially in the early stages of exploration when data are scarce, the skewness signal tends to be weak and significantly affected by noise. As a consequence, the direction selected by the biasing walls—based on this unreliable signal—amounts to an effectively random search (see Figure S5, Page 10, see SI). The bias then fails to consistently steer the system along meaningful reaction paths, greatly reducing the method's effectiveness.

These limitations motivated the design of the Skewencoder architecture, where the latent space is explicitly optimized to enhance the skewness signal, thereby providing a more robust and directional driving force for exploration.

Reviewer: Thanks for this additional investigation. I believe that these results should be also mentioned in the main text, since they clarify the need to use the proposed architecture as opposed to a plain autoencoder.

#Response#:
We thank the reviewer for this suggestion. We have added the following sentences to Section VI.A of the main text to direct the reader to the corresponding discussion in the SI.

"The effectiveness and efficiency of incorporating $\mathcal{L}_{skew}$ into the loss function are quantitatively analyzed using benchmark model potential (see SI)."

- Minimum Free Energy Pathways: How can it be guaranteed that the discovered paths correspond to minimum free energy paths? Except for the first toy model, the others only have one pathway. The last example has two pathways, and the fact that it can also find a higher pathway seems to contradict this aspect. Also, it seems that this will depend on the hyperparameters (e.g., κ, ϵ)? Is prior knowledge of energy barriers required to set them?

#Response#:
We appreciate the reviewer's comment and we gladly introduced further information in the submitted revised version.

Firstly, we must clarify that κ is not a hyperparameter of the Skewencoder but rather a variable for the exploration process. As indicated in the revised Figure S1 in SI (Page 4, see SI) (also referenced in Reviewer #1's Point 6), no prior knowledge of κ is necessary. In the newly updated version of Section IV.C (Page 11, Line 712-731, see main text), we specify that instead of sequentially increasing κ after each exploration campaign, our method deploys a small swarm of parallel explorations, each with distinct κ values. This approach facilitates more efficient sampling of the reaction network while maintaining control over computational cost and time. Thus, κ is initialized as a parameter for agnostic exploration of Loxodynamics to thoroughly and efficiently search the chemical reaction space.

Concerning the influence of κ, we invite the reviewer to refer to Point 7 from Reviewer #1, where we detailed our evaluation of exploration performance by initializing various κ values within the Loxodynamics simulation applied to the double-well model PES illustrated in the main text. Briefly, a higher κ tends to identify reaction pathways with elevated reaction barriers, whereas a suitably lower κ is inclined to discover pathways with reduced reaction barriers, closely approximating the anticipated MEP. Thus, selecting κ can significantly affect the exploration of distinct reaction pathways.

In the original manuscript, we had already discussed the role of both σ and ϵ in determining the position of the harmonic wall within the latent space (Page 10, Line 699-705, see main text). In the revised version, we have expanded the methodology section and included a new figure (Figure 7, Page 10, see main text). This figure illustrates that ϵ acts as a tunable offset, effectively controlling the degree of "brute force" applied during Loxodynamics simulations to explore new regions of the landscape.

Importantly, by adjusting the newly-added wall based on projected sample variance, this wall can explore unknown chemical spaces regardless of ϵ's value. Provided that ϵ is greater than 0, it will not constrain the range of extended chemical space exploration. Therefore, Loxodynamics' performance is not particularly sensitive to ϵ's selection, especially when it is small and positive. Since all sampling data are projected into the latent space and normalized—i.e., explored (known) space latent CVs are normalized to [-1, 1]—selecting ϵ at 0.1 represents a balanced choice that enhances simulation efficiency while preserving continuous and smooth exploration processes. The chosen ϵ remains consistent across all test cases, ensuring its neutrality concerning our method. In other words, controlling insensitive variables enables us to concentrate on analyzing κ's effect on Loxodynamics simulation performance.

Additionally, to ascertain whether the MD trajectory generated from Loxodynamics simulation approximates the MEP of the identified reaction pathway, we utilized the Climbing-Image Nudged Elastic Band (CINEB) method to optimize our MD trajectory snippet depicting reactions in both concerted and stepwise mechanisms (Section S2.C.3, Page 18, see SI). In summary, the geometries of the MD trajectory from Loxodynamics simulation closely resemble optimized configurational geometries on NEB paths, suggesting that the MD trajectory is already near the MEP of discovered reactions and indicating that Loxodynamics can guide systems approximately towards calculated MEP directions for certain reactions. Further simulation details are available in Section S2.C.3 (Page 18, see SI).

- The discussion section claims a wide range of potential applications, but the validation is limited to a few relatively simple cases (except for the last one). A more realistic assessment of current limitations, along with suggestions for how to transition from exploration to detailed mechanistic modeling, would strengthen the manuscript. For instance: after discovery, can the CVs produced by Skewencoder be used in conventional biased simulations to compute free energy surfaces?

#Response#:
We appreciate the reviewer for this insightful comment. In the newly added Section S2.B.3 (Page 15-16, see SI), we present results from well-tempered Metadynamics simulations employing the learned Skewencoder CV in both S_N2 and Diels-Alder reactions. These Skewencoder CVs are derived from iterations of Loxodynamics simulations where state transitions occurred. By properly configuring Metadynamics parameters, Figure S11 (Page 16, see SI) demonstrates that the Skewencoder CV can facilitate transitions between two states in Metadynamics simulations of both reactions. The Skewencoder CV also effectively differentiates reactant and product states, corresponding to the state information obtainable from heuristic CVs. Figure S12 (Page 16, see SI) illustrates that samples from well-tempered Metadynamics using Skewencoder CVs can reconstruct the FES for specific reactions. The reconstructed FES aligns with those generated by Metadynamics simulations employing heuristic CVs. Consequently, Skewencoder CVs can be further applied in conventional biased simulations, provided they encapsulate configurational information within both reactant and product spaces.

Moreover, we have included a new test case involving a significantly more complex chemical reaction network, which has been explored in an agnostic manner using Loxodynamics. Please check our answer to comment 4 to Reviewer #4.

Reviewer: My earlier comment referred specifically to the Discussion section, which still appears unchanged. Even with the additional results now included, there remains a noticeable gap between the scope of the presented demonstrations and the extensive list of potential applications described:

"... heterogeneous catalysis , reactive surface rearrangements , electrocatalysis , enzymatic catalysis , and asymmetric catalysis. Furthermore, it may offer insights into fundamental scientific questions such as prebiotic chemistry —including the Urey–Miller reactions—as well as environmentally and technologically relevant processes, including atmospheric and combustion chemistry, and pyrolysis . Likewise, it can be applied even to non-chemical activated events such as ligand–protein unbinding and polymorphic transitions in crystals, where new phases are not known a priori."

I would suggest making this section more balanced and realistic. It would also be helpful to explicitly acknowledge the current methodological constraints, for example, the need of system-specific restraining walls or state detections, and how these could be improved in future work.

#Response#:
We thank the reviewer for this suggestion. Accordingly, we have revised the corresponding part of the discussion as follows:

"This framework has the potential to drive significant advances in areas where the exploration of activated events is central, particularly in condensed-phase phenomena. These include surface-mediated processes such as heterogeneous catalysis—thermal, as exemplified by the cases presented here, electrochemical, and potentially photocatalytic—as well as selectivity-driven challenges in asymmetric, enzymatic, and supramolecular catalysis. Furthermore, the method may provide insights into other complex systems with unknown pathways, such as those found in prebiotic, atmospheric, and combustion chemistry. Conceptually, it is also applicable to non-chemical transformations like ligand–protein

unbinding, protein folding, or crystal nucleation and polymorphism, offering a tool to probe transitions where the final state is not known a priori."

We also explicitly acknowledge the necessity of system-specific walls and state detection methods in the Discussion section, as follows:

"We acknowledge that the current implementation relies on system-specific setups, including manually defined restraining walls and state-detection criteria, which still require limited but necessary chemical intuition and human intervention. Future work will focus on further automating these aspects, for instance through the development of an autonomous agent capable of managing the exploration phase, launching multiple exploration runs, designing and adapting biasing potentials on the fly, and automatically detecting new states while storing and organizing the associated information."

≫ **Methodological Details**

- It is unclear whether κ is chosen by the user or adapted during the iterations (as suggested. In the SN2 example, it is said: "Due to the asymmetric energy barriers… a larger harmonic constant was applied … to account for the higher barrier." Does setting this parameter require prior barrier knowledge?

#Response#:

We thank the reviewer for raising this point. In practice, we initiate a small swarm of independent trajectories (typically 5 to 10 depending on the complexity of the problem), each characterized by a different value of the biasing parameter κ. This strategy allows us to explore multiple directions efficiently and identify suitable paths for further exploration.

We acknowledge that this detail was not clearly explained in our original submission. To address this, we have updated both the main manuscript (Section IV.C, Page 11, see main text) and the Supporting Information (Figure S1, Pages 4, see SI), where we now explicitly describe this initial swarm approach. Additionally, we revised the flowchart in Figure S1 to reflect this procedure, clearly indicating the use of an initial set of test trajectories that guide the subsequent path exploration.

In general, no precise prior knowledge of the barrier is required.

Reviewer: While this answer is clear, I feel that this point is still not properly explained in the main text. In the general workflow section, it is described as: "The procedure is initialized with a sequence κ_i, where each value is distributed at uniform intervals…" I would suggest explaining more in detail like in the answer provided here instead ("We initiate a small swarm of independent trajectories … for further exploration.")

#Response#:

We thank the reviewer for this observation. We have amended the relevant text in Section VI.C to address this point, as shown below:

"The procedure is initialized with a small swarm of independent trajectories, each characterized by a different value of the biasing parameter κ_i. This strategy allows us to explore multiple directions efficiently and identify suitable paths for further exploration."

For the details of the Loxodynamics workflow, we have revised the corresponding text in the SI as follows:

"The procedure is initialized with a sequence $\{\kappa_i\}$, where each value is distributed at uniform intervals such that $\kappa_{i+1} = \kappa_i + \delta\kappa$, with $\delta\kappa$ denoting a constant increment."

- The rationale behind the inclusion of σ in the wall position (i.e., $\mu + \sigma + \epsilon$) is not explained in the main text. Figure 1 should also reflect this logic explicitly. Also, the role of ϵ seems important for determining the degree of exploration, but is not discussed. A discussion of how to choose these parameters and what effect they have is needed.

#Response#:
We appreciate the reviewer for highlighting this issue. In the newly revised manuscript, we provide a detailed explanation in the main text regarding the impact of utilizing projected sample variance σ and a custom offset ϵ, as illustrated in Figure 7 (Page 10, see main text), rather than merely portraying them in Figure 1 as before.

In the current revised version, Figure 1 (Page 3, see main text) primarily illustrates that the relationship between sample skewness and local anharmonicity arises from biased simulations on an artificial 1-D PES. During each iteration of the 1-D simulation, a bias harmonic wall is introduced and consistently shifted with a constant offset Δx from the left to the right along the x-axis. We observed that as the system transitioned from one minimum to another, the sample distribution became increasingly asymmetric, reflecting its gradual evolution towards state transition. In the current manuscript version, we have clarified that this observation from the 1-D simulation served as inspiration for going from a simple ideal case to more realistic scenarios. Consequently, to address challenges in extending this knowledge to multidimensional scenarios, we developed the Skewencoder NN architecture and Loxodynamics workflow as innovative tools for chemical space exploration.

The impact of the custom offset ϵ has been addressed in the response to the Method Validation section of Major Comments from Reviewer #2. Similarly, we employ projected sample variance to determine the harmonic wall's position. Calculating projected sample skewness necessitates first computing sample variance; thus, incorporating this term into our algorithm incurs no additional computational cost. Furthermore, sample variance reflects the distribution width within latent space. By using this term as an offset for the harmonic wall, previously explored regions in latent space are bypassed in new simulation iterations, reducing overlap and enhancing exploration efficiency. In practice, one may replace projected sample variance σ or even σ combined with ϵ with a constant offset (as illustrated in Figure 1, page 3, see main text). This substitution yields varying effects across different PES. By introducing σ and a moderate ϵ, we devised a mechanism that adaptively adjusts Loxodynamics' exploration "speed" based on local datasets without incurring extra computational expense.

Reviewer: I appreciate the additions, which I believe are helpful to better understand the methodology. However, I think that a few words on epsilon would still be useful, even just to give an idea of which values should be used based on your experience if, as it seems, it is not a very critical parameter. In addition, Fig.7 is not explained properly by the caption, in which many elements are not explained (colors, arrows, symbols, abbreviations..). In addition, it is not clear to me why the notation for the shift of the potential wall is different here (σ and ϵ) with respect to the notation used in Fig.1 (Δx).

#Response#:
We thank the reviewer for this further comment. We have expanded the discussion of ϵ in the main text, specifically by providing a range of values suitable for exploration. We believe this further elucidates the role of ϵ in Loxodynamics simulations. The revision is as follows:

"Values of $\epsilon \leq 1.0$ are considered moderate."

We apologize for the confusion regarding the symbol changes. We have revised the captions of both figures to clarify the notation. Specifically, we note that Δx in Figure 1 represents a constant shift within the biased sampling scheme that inspired Loxodynamics. In contrast, σ represents a statistical property of samples projected onto the latent space, while ϵ is a user-defined parameter. If one were to draw an analogy, the combined term $(\sigma + \epsilon)$ in Loxodynamics acts as the effective shift in the latent space—analogous to Δx—with the key distinction that $(\sigma + \epsilon)$ is dynamic rather than constant.

- The claim that the reaction discovery is "unbiased" is misleading. Since the skewness of the biased distribution depends on the applied parameters of the harmonic bias and no reweighting is performed, the method introduces a directional sampling artifact.

#Response#:
*We appreciate the reviewer for highlighting this issue. We would like to clarify that we do not claim the reaction discovery process to be **unbiased**, but rather **agnostic**. Although it is true that no reweighting is performed, such an action would primarily reproduce the equilibrium distribution of the initial state, providing limited insight into transition pathways. Therefore, a certain degree of directionality must be introduced. Importantly, this directional bias is not imposed a priori by the user but emerges during the early stages of the simulation when the system experiences minimal or no bias. As exploration progresses, the method selects specific directions for sampling reaction pathways. This selection is not manually adjusted but driven by an iterative sampling protocol that dynamically adapts to the underlying free energy landscape. For clarity, we have removed "unbiased" from the sentence: By following pathways with the lowest energy barriers, this approach enables efficient and unbiased exploration of likely reaction channels (Page 2, Line 101, see main text).*

- A complex procedure to decide whether to warm-start the autoencoder or not is put forward, but not justified. A comparison with an exploration training only from scratch on the toy model would clarify its impact. Also, this procedure requires a criterion to determine whether a new metastable state has been reached, which is not described for the examples reported. This contradicts the claim of minimal user

intervention and should be detailed more clearly.

#Response#:
We appreciate the reviewer for raising this insightful issue.

In the revised version of our manuscript, we have included Section S1.C (Page 5, see SI) to elucidate the criterion for distinguishing between metastable states and non-metastable states. Generally, this criterion can be adapted to any connectivity matrix-based method capable of determining atomic connectivity through customized switching functions. The present implementation is an in-house simplified version designed for efficiency.

Additionally, a comparative experiment was conducted between the Loxodynamics simulation utilizing the warm start training strategy and exploration training initiated from scratch on the toy model PES. The results are presented in Section S2.A.3 (Page 11, see SI). In summary, the warm start training strategy enhances exploration efficiency by lowering training and computational costs, thereby fulfilling the intended objective of this strategy within Loxodynamics simulation.

Reviewer: While I appreciate the additional material provided in the Supporting Information, I still do not understand why these elements are not at least mentioned in the main text. They appear highly relevant, particularly in light of the results shown in Figure S7. A brief reference or summary in the main text would help readers better grasp their significance.

#Response#:
We appreciate this suggestion. We have added a brief summary to Section VI.C of the main text as follows:

"To demonstrate the enhanced efficiency of the warm start strategy, we compared Loxodynamics simulations utilizing this approach against exploration training initiated from scratch (see SI)."

- Section A of the methods could benefit from a few revisions. I think that the workings of neural network layers are now part of common knowledge and so can be moved to the SI.

#Response#:
We thank the reviewer for this valuable suggestion. This content has been relocated to Section S1.B.2 (Page 3, Line 43-51, in SI), integrated with the discussion on the NN structure of Skewencoder, in the newly revised version of the manuscript.

On the other hand, it is not clear if the details about the $\mathcal{L}_{\mathtt{skew}}$ are just conceptual considerations (as it seems to me looking at the code) or technical details that need to be implemented to improve the performance of the method. Whichever the case, I think this should be made clear to the reader.

#Response#:

We appreciate the reviewer for identifying this issue. The current format of $\mathcal{L}_{skew}$ encompasses both conceptual considerations and technical implementation, enhancing the method's performance. Refer to Point 4 from Reviewer #1 and the discussion in the newly added Section S2.A.1 (Page 8-9, in SI) regarding this matter. In summary, the current implementation of $\mathcal{L}_{skew}$ has been shown to improve exploration efficiency when compared with other forms of $\mathcal{L}_{skew}$. Additionally, Skewencoder utilizing $\mathcal{L}_{skew}$ in a purely Gaussian-like form successfully identified alternative states and exhibited only slightly reduced performance during Loxodynamics simulation on the toy PES, suggesting that $\mathcal{L}_{skew}$ can be applied flexibly and is therefore also a conceptual consideration.

- In general, it would be much better to include in the methods section a summary of the method, including the loss functions (which are only defined in the introduction) and the workflow of the iterative procedure (which is again discussed in the introduction and summarized only in the SI).

#Response#:
We appreciate the reviewer for this helpful suggestion. We have reorganized the methodology section in the main text accordingly. To eliminate redundancy, we continue to focus on discussing the design of the skewness loss-related term in Section IV.A (Page 9-10, see main text), but have removed the introduction to a general NN structure. Subsequently, we partially revised Section VI.B (Page 10, see main text). We included Figure 7 (Page 10, see main text), which more clearly illustrates how the harmonic wall is applied based on sample mean, variance, and skewness values. Additionally, we employed a simplified equation form indicating how the harmonic wall is applied in latent space, as suggested by Reviewer #1 in Point 3. As a conclusion to the methodological section, we present a simplified version of the entire algorithm in Section IV.C (Page 11, see main text), corresponding to detailed discussions in Section S1 of the SI.

Comments specific to the applications presented.

$\gg$ **Toy Model**
- The evolution of the skewness and the mean of the CV across iterations is not shown. Including this would help illustrate how the method progresses and adapts the biasing direction based on the measured skewness?

#Response#:
We appreciate the reviewer for highlighting this issue. We wish to clarify that the latent space undergoes changes at each iteration, rendering comparisons of distribution moments across different variables impractical.

Concerning CV evolution as requested by Reviewer #1, please refer to Point 1 from Reviewer #1 where we illustrate the evolution of CVs in toy models.

- How does the exploration change as a function of the chosen hyperparameters, κ and ϵ? Showing supplementary examples would help clarify how to set them.

#Response#:
We appreciate the reviewer for this comment. We have previously addressed the impact of ϵ in earlier questions (e.g., Method Validation section of Major Comments from Reviewer #2). Regarding the effect of κ, as discussed in previous feedback, briefly, a higher κ tends to identify reaction pathways with elevated reaction barriers, whereas a suitably lower

κ is inclined to discover pathways with reduced reaction barriers, closely approximating the anticipated MEP. Thus, selecting κ can significantly influence the exploration of distinct reaction pathways. We have added Section S2.A.2 (Page 9-11, see SI) to assess the effect of κ using a toy model in Loxodynamics.

Concerning Skewencoder hyperparameters, we assert that with a typical deep autoencoder neural network structure, Skewencoder can effectively learn features from latent space and drive Loxodynamics. We have already compared Skewencoder performance using both deeper and shallower neural network structures in Section S2.B.2 (Page 13-15, see SI). Both configurations of the neural network successfully drive the system to identify state transitions. Therefore, given the robust fitting capability of neural networks, a general Autoencoder structure can be employed within Skewencoder without extensive hyperparameter optimization efforts. This highlights our method's flexibility and scalability.

- Although the potential includes two distinct paths, only one is sampled. Can the method recover both, depending on the hyperparameters used? This would also help clarify how this is obtained in the chabazite case.

#Response#:
We appreciate the reviewer for highlighting this issue. We kindly request the reviewer to refer to Point 7 from Reviewer #1 and Section S2.A.2 (Page 9-11, see SI).

$\gg$ **SN2 Reaction**
- The input CV space essentially consists of d_{Cl} and d_F, which already capture the transition without needing any combination of them. In my opinion, this system should be moved to the SI, as it does not significantly validate the method.

#Response#:
We appreciate the reviewer for this suggestion. However, we would like to emphasize that although the S_N2 reaction is a relatively simple example, it serves as a realistic test case. Given that the input dimensions of the previous model PES are all 2D, and the S_N2 reaction involves two Skewencoder input descriptors, it can be considered an appropriate intermediary between more complex real-world examples and simpler model PES, which are easier to test. Consequently, we have decided to retain this section in the main text, as it continues to validate the method.

$\gg$ **Chabazite**
- How is the second pathway found? In the SI it is said that different elastic constants are used to reflect barrier differences. Since barrier information is usually unknown, how would one proceed in general?

#Response#:
This is certainly a fundamental point. As discussed in our response to Point 6 from Reviewer #1 and illustrated in Figure S1 (Page 4, see SI), our method initiates a small swarm of parallel explorations, each with a distinct κ value that increases within a specified range. This approach allows Loxodynamics simulation to identify corresponding reaction pathways with different κ values without requiring barrier information. In the SI, we mention that various constants are employed to reflect barrier differences, indicating that different κ values configuring harmonic walls correspond to barrier height differences calculated using DFT methods (as reported in https://pubs.acs.org/doi/10.

1021/jp513024z). We do not imply that κ directly corresponds to energy values; rather, smaller κ values are likely to uncover reaction pathways with lower barriers, whereas larger κ values tend to reveal pathways with higher barriers. This observation aligns with our findings in the double pathway model PES presented in Section S2.A.2 (Page 9-11, see SI).

- In the stepwise mechanism, how is backward motion prevented? Are additional walls required? Would this become limiting for more complex catalytic cycles involving many intermediate steps?

#Response#:
We appreciate this comment very much as we realize this point was not clear in the original submission. Thus, we updated accordingly Section S2.C.2(Page 18, see SI), discussing how we restart the Loxodynamics simulation to explore a reaction pathway that includes multiple reaction wells, e.g. the stepwise mechanism in the ethanol dehydration case. In brief, Figure S13 (Page 17, see SI) shows that once we reached the intermediate state, we restart the Loxodynamcis, keep all the parameters related to the Skewencoder, but we retain the old CV along which we apply an additional restraining wall that guide the system to identify the current intermediate state. Then in the newly restarted Loxodynamics simulation, an "unbiased" simulation will be conducted, with the restraining wall kept. Then new Loxodynamics bias will be added in the following iteration, with the restraining wall still existing. To compensate the additional potential bias caused by the restraining wall, κ value must be increased to explore the second potential well, though the barrier height of the first reaction and the second reaction in the multiple-well reaction path are similar. We reported this observation in Section S2.C.2 (Page 18, see SI) and S2.D (Page 18-20, see SI), corresponding to different dehydration reactions on the acidic Chabasite. Based on our observation, we assume that this might be a limitation to a reaction pathway with more wells. However, in the test cases we have conducted Loxodynamics simulations on so far we do not observe such a limitation.

Reviewer: Also in this case, the use of a wall from the previous iteration is not described well in the text, despite sounding important in preventing back-crossing. The detailed explanation given here should thus somehow be incorporated in the main text. I also believe that Fig. S13 should at least be mentioned in the main text contextually, if not added directly to the main text.

#Response#:
We thank the reviewer for this suggestion. We have added a brief discussion to the main text regarding how back-crossing is prevented after restarting a Loxodynamics simulation. The revised passage is as follows:

"To prevent back-crossing to the original state upon reaching this intermediate state, we restart the Loxodynamics simulation. While the Skewencoder parameters and the original CV are retained, an additional restraining wall is applied along this CV to confine the system to the current intermediate state. The simulation then proceeds with this constraint active. (For further details regarding the restart protocol, refer to the SI.)"

≫ **Minor Comments**
- Many figure captions are not self-contained. For example, the meaning of color schemes and symbols (e.g., Figure 1 symbols, Figure 3 colorbar, color of scatter plots, Figure 4 color of scatter plots, Figure S1

symbols, Figure S2 colorbars, Figure S4 color of scatter plots, colorbars) is only defined in the main text and should be reported in the captions.

#Response#:
We appreciate the reviewer for identifying this issue. Below, we detail the content added to the captions of the relevant figures to address this concern. Generally, we emphasize that the precise values on the colorbar (Figure 3 and Figure S2) for both the model PES and projected latent space (Figure S4) are not critical. The closed format has already been provided for the PES. As Loxodynamics simulations progress, the projected CV space evolves continuously, with exact values in projected latent space CV being updated iteratively. Consequently, relative relationships are more significant than exact values for clearly illustrating CV evolution.

In the caption of Figure 1 (Page 3, see main text), we added:

- The initial wall is located at $x = \mu$, corresponding to the mean of the unbiased sample distribution (indicated by red dots in the upper panel). In each subsequent iteration, the harmonic wall is translated by a fixed offset Δx relative to its position in the previous iteration. The spring constant κ for all walls (depicted as blue, green, and yellow half-parabolic curves in the lower panel) remains constant throughout all iterations.

Reviewer: As mentioned in a previous comment, the Authors should ensure that the notation used here is consistent with the one used in Fig. 7 (Δx vs ϵ and σ).

#Response#:
We thank the reviewer for raising this point. We have revised the captions of both figures to clarify our explanation.

In the caption of Figure 3 (Page 5, see main text), we added:

- The darker region signifies an area with lower potential values, whereas the lighter region denotes an area with higher potential values.

Reviewer: I would at least indicate the spacing of the isolines in the caption (same for Figure S3 below) to provide a quantitative information.

#Response#:
We thank the reviewer for this suggestion. We have added a color bar to Figure 3, which also aligns with the editor's requirements. For Figure S3, we have explicitly indicated the spacing of the isolines in the figure caption.

- The central and lower panels display the evolution of the x (in dark green) and y (in light green) coordinates, respectively, over time (in Lennard-Jones units). The blue points represent samples from the incremental dataset, while red points correspond to samples from the current iteration.

In the caption of Figure 4 (Page 6, see main text), we added:

- In the 2D maps, as in Figure 3, blue points denote samples from the global dataset, whereas red points indicate samples from the local dataset corresponding to the current iteration. Relevant structures extracted from the trajectory blocks of each successive iteration are shown in the upper subpanels.

In the caption of Figure S1(Page 4, see SI), we added:

- In this flowchart, $\{\kappa_i\}$ represents a sequence of wall spring constants arranged in an arithmetic progression with a common difference of $\delta\kappa$. The parameter n_{iter} denotes the maximum number of iterations permitted for the Loxodynamics simulation. The quantity γ_{ls} refers to the projected sample skewness in latent space. The variable $\mathbf{L}_{state}$ indicates the list of identified metastable states, while S_p serves as the identifier for the current state.

In the caption of Figure S3 (Page 7, see SI), we added:

- The darker region signifies an area with lower potential values, whereas the lighter region denotes an area with higher potential values.

In the caption of Figure S4 (Page 8, see SI), we added:

- In each left subcolumn, The blue points represent samples from the incremental dataset, while red points correspond to samples from the current iteration.

- Lighter regions correspond to higher values in the latent CV space, whereas darker regions represent lower values.

- The provided code is well-structured and detailed, which ensures reproducibility. However, from the SI and the code, it appears to rely extensively on the mlcolvar library, which is not cited in the main text.

#Response#:
We appreciate the reviewer for this suggestion. The relevant citation has been incorporated in the code and data availability section (Page 12, see main text).

- When quantities are discussed in the main text, refer to the specific equation number to help guide the reader (e.g., $\mathcal{L}_{skew}$ in methods sec. A).

#Response#:
We appreciate the reviewer for this suggestion. We have added the following in the main text (Page 9-11, line 631-633): This section is devoted to the discussion of the formulation and implementation of the skewness-related loss term, $\mathcal{L}_{skew}$ (defined in Equation 4), to address this issue.

Remarks on code availability:

The provided code is well-structured and detailed, which ensures reproducibility. However, from the SI and the code, it appears to rely extensively on the mlcolvar library, which is not cited in the main text.

Reviewer #3

Remarks to the Author:

I co-reviewed this manuscript with one of the reviewers who provided the listed reports. This is part of the Nature Communications initiative to facilitate training in peer review and to provide appropriate recognition for Early Career Researchers who co-review manuscripts.

Reviewer #4

Remarks to the Author:

I appreciate the effort that the authors have undertaken to reply to all the reviewers' comments. I find the additional examples and information included in the manuscript and SI to drastically improve the reader's understanding of the method as well as its relevance. Therefore, I would be happy to recommend the manuscript for publication in Nature Communications. I only have one final remark for the authors: in section S3, a comparison between the free energy profiles obtained from a chemically-inspired CV and the CV from Loxodynamics is reported (Figure S12). However, the free energy profile is CV dependent and comparing two profiles directly is not per se meaningful (see, e.g., https://doi.org/10.1016/j.jcat.2020.04.015 or https://doi.org/10.1063/5.0083423). I would find it interesting to evaluate the actual CV-independent free energy difference between the two metastable states from the simulations results, to showcase that the Loxodynamics CV is (likely) fine and discrepancies in the free energy profile with chemically-inspired CVs are meaningless.

#Response#:
We thank the reviewer for highlighting this point. We calculated the CV-independent free energy difference between the reactants and products for the SN2 and Diels-Alder reactions by integrating the probability densities within the respective basins. The free energy differences obtained from Loxodynamics sampling matched closely with those obtained using chemically-inspired CVs in these two test cases. Further details have been added to the SI, as revised below:

"The free energy differences between reactants and products were calculated by integrating the probability densities within the respective basins. The dividing boundary between basins was set to −0.28 *for the chemically-inspired CV and* 0.0 *for the Skewencoder CV. The chemically-inspired CV simulation yielded a free energy difference of* 143.177 *kJ/mol, while the Skewencoder CV simulation resulted in* 143.204 *kJ/mol, demonstrating excellent agreement between the two simulation results. ... Still, we calculated the free energy difference by integrating over basins defined by a threshold of* −2.2 *(chemically-inspired CV) and* −0.7 *(Skewencoder CV). The resulting values were* 106.223 *kJ/mol and* 102.342 *kJ/mol, respectively, indicating reasonable agreement between the two simulation results."*

Reviewer #2

I appreciate the Authors' efforts in addressing the reviewers' questions, which have clarified several important aspects of the work. That said, I believe that some minor aspects could still be refined to further improve the manuscript.

In particular, in several cases the clarifications provided in the response letter have not been incorporated into the manuscript, or have been included only in the Supporting Information without being mentioned in the main text. This makes it difficult for readers to access or appreciate these details. Since some of these issues were raised independently by multiple reviewers, adding the corresponding explanations or cross-references in the manuscript would help make the work clearer and more robust.

Furthermore, the inclusion of the more complex example (1-butanol dehydration) is a welcome improvement and represents a valuable extension of the study. However, this example lacks sufficient supporting data: In addition to the schematic of the "discovered" species, some quantitative evidence (e.g., time-series data or scatter plots) should be reported in the SI.

Regarding Section D.2, it would also be useful to briefly mention alternative or complementary approaches for state identification, as suggested by other reviewers; for instance, the use of machine-learning-based classifiers that could offer a more general or automated strategy.

Please find below a few additional points related to the original remarks.

Original remarks to the authors

The manuscript by Zhang and Piccini presents an original and promising approach to reaction discovery based on the skewness of probability distribution using an autoencoder framework. The idea behind the method—Loxodynamics—is clearly explained and introduces a creative use of probability distribution 4 asymmetry to guide sampling. However, several methodological and validation aspects require clarification and improvement before the work can be recommended for publication.

Major Comments

Related Work and Literature.

The paper lacks a thorough discussion of prior literature on reaction discovery using machine learning collective variables, especially autoencoder-based methods. Relevant prior work includes: We appreciate the reviewer's careful attention in highlighting fundamental literature that should be cited. Below, we outline the additions we have incorporated into the manuscript to address these suggestions.

- Chen & Ferguson https://doi.org/10.1063/1.5023804 demonstrated the use of autoencoders for interleaved CV discovery and enhanced sampling, iteratively expanding the sampled space.

#Response#:

We added the suggested reference to Chen and Ferguson J. Chem. Phys. one we already cited on the same topic (Page 3, Line 208-219, see main text).

- Ketkaew & Luber https://doi.org/10.1021/acs.jcim.2c00883 applied autoencoders to generate blind CVs for chemical reactions starting only from reactants.

#Response#: We added the suggested reference (Page 3, Line 219-231, see main text).

Reviewer: The way this reference is mentioned is not contextualized correctly; it seems to even imply that it already solved the problem addressed in the present work by the Authors, which contrasts with the clarification given below.

- Raucci et al. https://doi.org/10.1021/acs.jpclett.1c03993 proposed a method for blind reaction discovery using general-purpose CVs with OPES explore.

#Response#:
We added the suggested reference (Page 2, Line 90-91, see main text).
Reviewer: This reference is cited but even less contextualized. Considering that it addresses the same issue the Authors discuss, it should be at least mentioned in the part about similar approaches; together with the clarification below.

Nonetheless, we believe that a further clarification on the distinction between Loxodynamics and other schemes is warranted. The aforementioned approaches employ CVs together with enhanced sampling methods such as umbrella sampling, metadynamics, or OPES to accelerate sampling. These methods, however, do not encode a specific directionality. In contrast, Loxodynamics introduces a bias explicitly designed to drive the system toward a new state along the CV. Metadynamics and OPES, in particular, are "flooding" approaches, in which the bias gradually fills the local free-energy basin to enhance fluctuations within it. The philosophy of Loxodynamics is fundamentally different: rather than amplifying fluctuations, it reshapes them with the specific purpose of propelling the system uphill.

Reviewer: This answer clarifies the difference with respect to the other methods. Adding it to the manuscript would be highly beneficial.

Method Validation

Comparison with standard autoencoders: what advantage does the "Skewencoder" offer over a standard 5 autoencoder with the same architecture? A comparative experiment (e.g., on the toy model) would help clarify its contribution.

#Response#: We thank the reviewer for raising this important point. Indeed, in our initial attempts, we employed a standard Autoencoder to achieve our goal. This involved projecting the configurations onto the latent space and evaluating the skewness of the resulting distribution without explicitly optimizing for it. In principle, such an approach could work. Accordingly, we performed comparative experiments involving both the current Skewencoder configuration and a pure Autoencoder format (Section S2.A.1, Page 10-11, see SI), alongside evaluations of various formulations of the skewness loss term (refer to Reviewer #1's Point 4). However, we consistently observed that, especially in the early stages of exploration when data are scarce, the skewness signal tends to be weak and significantly affected by noise. As a consequence, the direction

selected by the biasing walls—based on this unreliable signal—amounts to an effectively random search (see Figure S5, Page 10, see SI). The bias then fails to consistently steer the system along meaningful reaction paths, greatly reducing the method's effectiveness. These limitations motivated the design of the Skewencoder architecture, where the latent space is explicitly optimized to enhance the skewness signal, thereby providing a more robust and directional driving force for exploration.

Reviewer: Thanks for this additional investigation. I believe that these results should be also mentioned in the main text, since they clarify the need to use the proposed architecture as opposed to a plain autoencoder.

Minimum Free Energy Pathways

How can it be guaranteed that the discovered paths correspond to minimum free energy paths? Except for the first toy model, the others only have one pathway. The last example has two pathways, and the fact that it can also find a higher pathway seems to contradict this aspect. Also, it seems that this will depend on the hyperparameters (e.g., κ, ϵ)? Is prior knowledge of energy barriers required to set them?

#Response#: We appreciate the reviewer's comment and we gladly introduced further information in the submitted revised version. Firstly, we must clarify that κ is not a hyperparameter of the Skewencoder but rather a variable for the exploration process. As indicated in the revised Figure S1 in SI (Page 4, see SI) (also referenced in Reviewer #1's Point 6), no prior knowledge of κ is necessary. In the newly updated version of Section IV.C (Page 11, Line 712-731, see main text), we specify that instead of sequentially increasing κ after each exploration campaign, our method deploys a small swarm of parallel explorations, each with distinct κ values. This approach facilitates more efficient sampling of the reaction network while maintaining control over computational cost and time. Thus, κ is initialized as a parameter for agnostic exploration of Loxodynamics to thoroughly and efficiently search the chemical reaction space. Concerning the influence of κ, we invite the reviewer to refer to Point 7 from Reviewer #1, where we detailed our evaluation of exploration performance by initializing various κ values within the Loxodynamics simulation applied to the double-well model PES illustrated in the main text. Briefly, a higher κ tends to identify reaction pathways with elevated reaction barriers, whereas a suitably lower κ is inclined to discover pathways with reduced reaction barriers, closely approximating the anticipated MEP. Thus, selecting κ can significantly affect the exploration of distinct reaction 6 pathways. In the original manuscript, we had already discussed the role of both σ and ϵ in determining the position of the harmonic wall within the latent space (Page 10, Line 699-705, see main text). In the revised version, we have expanded the methodology section and included a new figure (Figure 7, Page 10, see main text). This figure illustrates that ϵ acts as a tunable offset, effectively controlling the degree of "brute force" applied during Loxodynamics simulations to explore new regions of the landscape. Importantly, by adjusting the newly-added wall based on projected sample variance, this wall can explore unknown chemical spaces regardless of ϵ's value. Provided that ϵ is greater than 0, it will not constrain the range of extended chemical space exploration. Therefore, Loxodynamics' performance is not particularly sensitive to ϵ's selection, especially when it is small and positive. Since all sampling data are projected into the latent space and normalized—i.e., explored (known) space latent CVs are normalized to [-1, 1]—selecting ϵ at 0.1 represents a balanced choice that enhances simulation efficiency while preserving continuous and smooth exploration processes. The chosen ϵ remains consistent across all test cases, ensuring its neutrality concerning our method. In other words, controlling insensitive variables enables us to concentrate on analyzing κ's effect on Loxodynamics simulation performance. Additionally, to ascertain whether the MD trajectory generated from Loxodynamics simulation approximates the

MEP of the identified reaction pathway, we utilized the Climbing-Image Nudged Elastic Band (CINEB) method to optimize our MD trajectory snippet depicting reactions in both concerted and stepwise mechanisms (Section S2.C.3, Page 18, see SI). In summary, the geometries of the MD trajectory from Loxodynamics simulation closely resemble optimized configurational geometries on NEB paths, suggesting that the MD trajectory is already near the MEP of discovered reactions and indicating that Loxodynamics can guide systems approximately towards calculated MEP directions for certain reactions. Further simulation details are available in Section S2.C.3 (Page 18, see SI).

The discussion section claims a wide range of potential applications, but the validation is limited to a few relatively simple cases (except for the last one). A more realistic assessment of current limitations, along with suggestions for how to transition from exploration to detailed mechanistic modeling, would strengthen the manuscript. For instance: after discovery, can the CVs produced by Skewencoder be used in conventional biased simulations to compute free energy surfaces?

#Response#: We appreciate the reviewer for this insightful comment. In the newly added Section S2.B.3 (Page 15-16, see SI), we present results from well-tempered Metadynamics simulations employing the learned Skewencoder CV in both SN2 and Diels-Alder reactions. These Skewencoder CVs are derived from iterations of Loxodynamics simulations where state transitions occurred. By properly configuring Metadynamics parameters, Figure S11 (Page 16, see SI) demonstrates that the Skewencoder CV can facilitate transitions between two states in Metadynamics simulations of both reactions. The Skewencoder CV also effectively differentiates reactant and product states, corresponding to the state information obtainable from heuristic CVs. Figure S12 (Page 16, see SI) illustrates that samples from well-tempered Metadynamics using Skewencoder CVs can reconstruct the FES for specific reactions. The reconstructed FES aligns with those generated by Metadynamics simulations employing heuristic CVs. Consequently, Skewencoder CVs can be further applied in conventional biased simulations, provided they encapsulate configurational information within both reactant and product spaces. 7 Moreover, we have included a new test case involving a significantly more complex chemical reaction network, which has been explored in an agnostic manner using Loxodynamics. Please check our answer to comment 4 to Reviewer #4.

Reviewer: My earlier comment referred specifically to the Discussion section, which still appears unchanged. Even with the additional results now included, there remains a noticeable gap between the scope of the presented demonstrations and the extensive list of potential applications described:

"...heterogeneous catalysis, reactive surface rearrangements, electrocatalysis, enzymatic catalysis, and asymmetric catalysis. Furthermore, it may offer insights into fundamental scientific questions such as prebiotic chemistry—including the Urey–Miller reactions—as well as environmentally and technologically relevant processes, including atmospheric and combustion chemistry, and pyrolysis. Likewise, it can be applied even to non-chemical activated events such as ligand–protein unbinding and polymorphic transitions in crystals, where new phases are not known a priori."

I would suggest making this section more balanced and realistic. It would also be helpful to explicitly acknowledge the current methodological constraints, for example, the need of system-specific restraining walls or state detections, and how these could be improved in future work.

Methodological Details

It is unclear whether κ is chosen by the user or adapted during the iterations (as suggested. In the SN2 example, it is said: "Due to the asymmetric energy barriers. . . a larger harmonic constant was applied . . . to account for the higher barrier." Does setting this parameter require prior barrier knowledge?

#Response#: We thank the reviewer for raising this point. In practice, we initiate a small swarm of independent trajectories (typically 5 to 10 depending on the complexity of the problem), each characterized by a different value of the biasing parameter κ. This strategy allows us to explore multiple directions efficiently and identify suitable paths for further exploration. We acknowledge that this detail was not clearly explained in our original submission. To address this, we have updated both the main manuscript (Section IV.C, Page 11, see main text) and the Supporting Information (Figure S1, Pages 4, see SI), where we now explicitly describe this initial swarm approach. Additionally, we revised the flowchart in Figure S1 to reflect this procedure, clearly indicating the use of an initial set of test trajectories that guide the subsequent path exploration. In general, no precise prior knowledge of the barrier is required.

Reviewer: While this answer is clear, I feel that this point is still not properly explained in the main text. In the general workflow section, it is described as: "*The procedure is initialized with a sequence {κi}, where each value is distributed at uniform intervals…*" I would suggest explaining more in detail like in the answer provided here instead *("We initiate a small swarm of independent trajectories … for further exploration.")*

The rationale behind the inclusion of σ in the wall position (i.e., μ + σ + ϵ) is not explained in the main text. Figure 1 should also reflect this logic explicitly. Also, the role of ϵ seems important for determining the degree of exploration, but is not discussed. A discussion of how to choose these parameters and what effect they have is needed.

#Response#
We appreciate the reviewer for highlighting this issue. In the newly revised manuscript, we provide a detailed explanation in the main text regarding the impact of utilizing projected sample variance σ and a custom offset ϵ, as illustrated in Figure 7 (Page 10, see main text), rather than merely portraying them in Figure 1 as before. In the current revised version, Figure 1 (Page 3, see main text) primarily illustrates that the relationship between sample skewness and local anharmonicity arises from biased simulations on an artificial 1-D PES. During each iteration of the 1-D simulation, a bias harmonic wall is introduced and consistently shifted with a constant offset Δx from the left to the right along the x-axis. We observed that as the system transitioned from one minimum to another, the sample distribution became increasingly asymmetric, reflecting its gradual evolution towards state transition. In the current manuscript version, we have clarified that this observation from the 1-D simulation served as inspiration for going from a simple ideal case to more realistic scenarios. Consequently, to address challenges in extending this knowledge to multidimensional scenarios, we developed the Skewencoder NN architecture and Loxodynamics workflow as innovative tools for chemical space exploration. The impact of the custom offset ϵ has been addressed in the response to the Method Validation section of Major Comments from Reviewer #2. Similarly, we employ projected sample variance to determine the harmonic wall's position. Calculating projected sample skewness necessitates first computing sample variance; thus, incorporating this term into our algorithm incurs no additional computational cost. Furthermore, sample variance reflects the distribution width within latent space. By using this term as an offset for the harmonic wall, previously explored regions in latent space are bypassed in new simulation iterations, reducing overlap and enhancing exploration efficiency. In practice, one may replace

projected sample variance σ or even σ combined with ϵ with a constant offset (as illustrated in Figure 1, page 3, see main text). This substitution yields varying effects across different PES. By introducing σ and a moderate ϵ, we devised a mechanism that adaptively adjusts Loxodynamics' exploration "speed" based on local datasets without incurring extra computational expense.

Reviewer: I appreciate the additions, which I believe are helpful to better understand the methodology. However, I think that a few words on epsilon would still be useful, even just to give an idea of which values should be used based on your experience if, as it seems, it is not a very critical parameter.
In addition, Fig.7 is not explained properly by the caption, in which many elements are not explained (colors, arrows, symbols, abbreviations..). In addition, it is not clear to me why the notation for the shift of the potential wall is different here (δ and ε) with respect to the notation used in Flg.1 (Δx).

The claim that the reaction discovery is "unbiased" is misleading. Since the skewness of the biased distribution depends on the applied parameters of the harmonic bias and no reweighting is performed, the method introduces a directional sampling artifact.

#Response#: We appreciate the reviewer for highlighting this issue. We would like to clarify that we do not claim the reaction discovery process to be unbiased, but rather agnostic. Although it is true that no reweighting is performed, such an action would primarily reproduce the equilibrium distribution of the initial state, providing limited insight into transition pathways. Therefore, a certain degree of directionality must be introduced. Importantly, this directional bias is not imposed a priori by the user but emerges during the early stages of the simulation when the system experiences minimal or no bias. As exploration progresses, the method selects specific directions for sampling reaction pathways. This selection is not manually adjusted but driven by an iterative sampling protocol that dynamically adapts to the underlying free energy landscape. For clarity, we have removed "unbiased" from the sentence: By following pathways with the lowest energy barriers, this approach enables efficient and unbiased exploration of likely reaction channels (Page 2, Line 101, see main text).

A complex procedure to decide whether to warm-start the autoencoder or not is put forward, but not justified. A comparison with an exploration training only from scratch on the toy model would clarify its impact. Also, this procedure requires a criterion to determine whether a new metastable state has been reached, which is not described for the examples reported. This contradicts the claim of minimal user intervention and should be detailed more clearly.

#Response#: We appreciate the reviewer for raising this insightful issue. In the revised version of our manuscript, we have included Section S1.C (Page 5, see SI) to elucidate the criterion for distinguishing between metastable states and non-metastable states. Generally, this criterion can be adapted to any connectivity matrix-based method capable of determining atomic connectivity through customized switching functions. The present implementation is an in-house simplified version designed for efficiency. Additionally, a comparative experiment was conducted between the Loxodynamics simulation utilizing the warm start training strategy and exploration training initiated from scratch on the toy model PES. The results are presented in Section S2.A.3 (Page 11, see SI). In summary, the warm start training strategy enhances exploration efficiency by lowering training and computational costs, thereby fulfilling the intended objective of this strategy within Loxodynamics simulation.

Reviewer: While I appreciate the additional material provided in the Supporting Information, I still do not understand why these elements are not at least mentioned in the main text. They appear highly relevant, particularly in light of the results shown in Figure S7. A brief reference or summary in the main text would help readers better grasp their significance.

Section A of the methods could benefit from a few revisions. I think that the workings of neural network layers are now part of common knowledge and so can be moved to the SI.

#Response#:
We thank the reviewer for this valuable suggestion. This content has been relocated to Section S1.B.2 (Page 3, Line 43-51, in SI), integrated with the discussion on the NN structure of Skewencoder, in the newly revised version of the manuscript.

On the other hand, it is not clear if the details about the Lskew are just conceptual considerations (as it seems to me looking at the code) or technical details that need to be implemented to improve the performance of the method. Whichever the case, I think this should be made clear to the reader.

#Response#: We appreciate the reviewer for identifying this issue. The current format of Lskew encompasses both conceptual considerations and technical implementation, enhancing the method's performance. Refer to Point 4 from Reviewer #1 and the discussion in the newly added Section S2.A.1 (Page 8-9, in SI) regarding this matter. In summary, the current implementation of Lskew has been shown to improve exploration efficiency when compared with other forms of Lskew. Additionally, Skewencoder utilizing Lskew in a purely Gaussian-like form successfully identified alternative states and exhibited only slightly reduced performance during Loxodynamics simulation on the toy PES, suggesting that Lskew can be applied flexibly and is therefore also a conceptual consideration.

In general, it would be much better to include in the methods section a summary of the method, including the loss functions (which are only defined in the introduction) and the workflow of the iterative procedure (which is again discussed in the introduction and summarized only in the SI).

#Response#: We appreciate the reviewer for this helpful suggestion. We have reorganized the methodology section in the main text accordingly. To eliminate redundancy, we continue to focus on discussing the design of the skewness loss-related term in Section IV.A (Page 9-10, see main text), but have removed the introduction to a general NN structure. Subsequently, we partially revised Section VI.B (Page 10, see main text). We included Figure 7 (Page 10, see main text), which more clearly illustrates how the harmonic wall is applied based on sample mean, variance, and skewness values. Additionally, we employed a simplified equation form indicating how the harmonic wall is applied in latent space, as suggested by Reviewer #1 in Point 3. As a conclusion to the methodological section, we present a simplified version of the entire algorithm in Section IV.C (Page 11, see main text), corresponding to detailed discussions in Section S1 of the SI.

Comments specific to the applications presented.

Toy Model

The evolution of the skewness and the mean of the CV across iterations is not shown. Including this would help illustrate how the method progresses and adapts the biasing direction based on the measured skewness?

#Response#:
We appreciate the reviewer for highlighting this issue. We wish to clarify that the latent space undergoes changes at each iteration, rendering comparisons of distribution moments across different variables impractical. Concerning CV evolution as requested by Reviewer #1, please refer to Point 1 from Reviewer #1 where we illustrate the evolution of CVs in toy models.

How does the exploration change as a function of the chosen hyperparameters, κ and ϵ? Showing supplementary examples would help clarify how to set them.

#Response#:
We appreciate the reviewer for this comment. We have previously addressed the impact of ϵ in earlier questions (e.g., Method Validation section of Major Comments from Reviewer #2). Regarding the effect of κ, as discussed in previous feedback, briefly, a higher κ tends to identify reaction pathways with elevated reaction barriers, whereas a suitably lower κ is inclined to discover pathways with reduced reaction barriers, closely approximating the anticipated MEP. Thus, selecting κ can significantly influence the exploration of distinct reaction pathways. We have added Section S2.A.2 (Page 9-11, see SI) to assess the effect of κ using a toy model in Loxodynamics. Concerning Skewencoder hyperparameters, we assert that with a typical deep autoencoder neural network structure, Skewencoder can effectively learn features from latent space and drive Loxodynamics. We have already compared Skewencoder performance using both deeper and shallower neural network structures in Section S2.B.2 (Page 13-15, see SI). Both configurations of the neural network successfully drive the system to identify state transitions. Therefore, given the robust fitting capability of neural networks, a general Autoencoder structure can be employed within Skewencoder without extensive hyperparameter optimization efforts. This highlights our method's flexibility and scalability.

Although the potential includes two distinct paths, only one is sampled. Can the method recover both, depending on the hyperparameters used? This would also help clarify how this is obtained in the chabazite case.

#Response#:
We appreciate the reviewer for highlighting this issue. We kindly request the reviewer to refer to Point 7 from Reviewer #1 and Section S2.A.2 (Page 9-11, see SI).

SN2 Reaction

The input CV space essentially consists of dCl and dF, which already capture the transition without needing any combination of them. In my opinion, this system should be moved to the SI, as it does not significantly validate the method.

#Response#:
We appreciate the reviewer for this suggestion. However, we would like to emphasize that although the SN2 reaction is a relatively simple example, it serves as a realistic test case. Given that the input dimensions of the previous model PES are all 2D, and the SN2 reaction involves two Skewencoder input descriptors, it can be considered an appropriate intermediary between more

complex real-world examples and simpler model PES, which are easier to test. Consequently, we have decided to retain this section in the main text, as it continues to validate the method.

Chabazite

How is the second pathway found? In the SI it is said that different elastic constants are used to reflect barrier differences. Since barrier information is usually unknown, how would one proceed in general?

#Response#: This is certainly a fundamental point. As discussed in our response to Point 6 from Reviewer #1 and illustrated in Figure S1 (Page 4, see SI), our method initiates a small swarm of parallel explorations, each with a distinct κ value that increases within a specified range. This approach allows Loxodynamics simulation to identify corresponding reaction pathways with different κ values without requiring barrier information. In the SI, we mention that various constants are employed to reflect barrier differences, indicating that different κ values configuring harmonic walls correspond to barrier height differences calculated using DFT methods (as reported in https://pubs.acs.org/doi/10. 1021/jp513024z). We do not imply that κ directly corresponds to energy values; rather, smaller κ values are likely to uncover reaction pathways with lower barriers, whereas larger κ values tend to reveal pathways with higher barriers. This observation aligns with our findings in the double pathway model PES presented in Section S2.A.2 (Page 9-11, see SI).

In the stepwise mechanism, how is backward motion prevented? Are additional walls required? Would this become limiting for more complex catalytic cycles involving many intermediate steps?

#Response#: We appreciate this comment very much as we realize this point was not clear in the original submission. Thus, we updated accordingly Section S2.C.2(Page 18, see SI), discussing how we restart the Loxodynamics simulation to explore a reaction pathway that includes multiple reaction wells, e.g. the stepwise mechanism in the ethanol dehydration case. In brief, Figure S13 (Page 17, see SI) shows that once we reached the intermediate state, we restart the Loxodynamcis, keep all the parameters related to the Skewencoder, but we retain the old CV along which we apply an additional restraining wall that guide the system to identify the current intermediate state. Then in the newly restarted Loxodynamics simulation, an “unbiased” simulation will be conducted, with the restraining wall kept. Then new Loxodynamics bias will be added in the following iteration, with the restraining wall still existing. To compensate the additional potential bias caused by the restraining wall, κ value must be increased to explore the second potential well, though the barrier height of the first reaction and the second reaction in the multiple-well reaction path are similar. We reported this observation in Section S2.C.2 (Page 18, see SI) and S2.D (Page 18-20, see SI), corresponding to different dehydration reactions on the acidic Chabasite. Based on our observation, we assume that this might be a limitation to a reaction pathway with more wells. However, in the test cases we have conducted Loxodynamics simulations on so far we do not observe such a limitation.

Reviewer: Also in this case, the use of a wall from the previous iteration is not described well in the text, despite sounding important in preventing back-crossing. The detailed explanation given here should thus somehow be incorporated in the main text. I also believe that Fig. S13 should at least be mentioned in the main text contextually, if not added directly to the main text.

Minor Comments

Many figure captions are not self-contained. For example, the meaning of color schemes and symbols (e.g., Figure 1 symbols, Figure 3 colorbar, color of scatter plots, Figure 4 color of scatter plots, Figure S1 symbols, Figure S2 colorbars, Figure S4 color of scatter plots, colorbars) is only defined in the main text 12 and should be reported in the captions.

#Response#: We appreciate the reviewer for identifying this issue. Below, we detail the content added to the captions of the relevant figures to address this concern. Generally, we emphasize that the precise values on the colorbar (Figure 3 and Figure S2) for both the model PES and projected latent space (Figure S4) are not critical. The closed format has already been provided for the PES. As Loxodynamics simulations progress, the projected CV space evolves continuously, with exact values in projected latent space CV being updated iteratively. Consequently, relative relationships are more significant than exact values for clearly illustrating CV evolution.

In the caption of Figure 1 (Page 3, see main text), we added:
- The initial wall is located at x = μ, corresponding to the mean of the unbiased sample distribution (indicated by red dots in the upper panel). In each subsequent iteration, the harmonic wall is translated by a fixed offset Δx relative to its position in the previous iteration. The spring constant κ for all walls (depicted as blue, green, and yellow half-parabolic curves in the lower panel) remains constant throughout all iterations.

Reviewer: As mentioned in a previous comment, the Authors should ensure that the notation used here is consistent with the one used in Fig. 7 (Δx vs δ and ε).

In the caption of Figure 3 (Page 5, see main text), we added:
- The darker region signifies an area with lower potential values, whereas the lighter region denotes an area with higher potential values.

Reviewer: I would at least indicate the spacing of the isolines in the caption (same for Figure S3 below) to provide a quantitative information.

- The central and lower panels display the evolution of the x (in dark green) and y (in light green) coordinates, respectively, over time (in Lennard-Jones units). The blue points represent samples from the incremental dataset, while red points correspond to samples from the current iteration.

In the caption of Figure 4 (Page 6, see main text), we added:
- In the 2D maps, as in Figure 3, blue points denote samples from the global dataset, whereas red points indicate samples from the local dataset corresponding to the current iteration. Relevant structures extracted from the trajectory blocks of each successive iteration are shown in the upper subpanels.

In the caption of Figure S1(Page 4, see SI), we added:
- In this flowchart, {κi} represents a sequence of wall spring constants arranged in an arithmetic progression with a common difference of δκ. The parameter niter denotes the maximum number of iterations permitted for the Loxodynamics simulation. The quantity γls refers to the projected sample skewness in latent space. The variable Lstate indicates the list of identified metastable states, while Sp serves as the identifier for the current state.

In the caption of Figure S3 (Page 7, see SI), we added:
- The darker region signifies an area with lower potential values, whereas the lighter region denotes an area with higher potential values.

In the caption of Figure S4 (Page 8, see SI), we added: 13
- In each left subcolumn, The blue points represent samples from the incremental dataset, while red points correspond to samples from the current iteration.
- Lighter regions correspond to higher values in the latent CV space, whereas darker regions represent lower values.

The provided code is well-structured and detailed, which ensures reproducibility. However, from the SI and the code, it appears to rely extensively on the mlcolvar library, which is not cited in the main text.

#Response#: We appreciate the reviewer for this suggestion. The relevant citation has been incorporated in the code and data availability section (Page 12, see main text).

When quantities are discussed in the main text, refer to the specific equation number to help guide the reader (e.g., Lskew in methods sec. A).

#Response#: We appreciate the reviewer for this suggestion. We have added the following in the main text (Page 9-11, line 631-633): This section is devoted to the discussion of the formulation and implementation of the skewness-related loss term, Lskew (defined in Equation 4), to address this issue.