

Path Sampling for Rare Events Boosted by Machine Learning

Commentary by

Porhouy Minh

Sapna Sarupria

Department of Chemistry and Chemical Theory Center, University of Minnesota Minneapolis, MN USA

on

Machine-guided path sampling to discover mechanisms of molecular self-organization

Hendrik Jung, Roberto Covino, A. Arjun, Christian Leitold, Christoph Dellago, Peter G. Bolhuis and Gerhard Hummer, [Nat. Comput. Sci. 3, 334–345 \(2023\)](#)

Received December 12, 2025

Published February 04, 2026

Statement of Significance

The study by Jung et al. [1] introduced Artificial Intelligence for Molecular Mechanism Discovery (AIMMD), a novel sampling algorithm that integrates machine learning to enhance the efficiency of transition path sampling (TPS). By enabling on-the-fly estimation of the committor probability and simultaneously deriving a human-interpretable reaction coordinate, AIMMD offers a robust framework for elucidating the mechanistic pathways of complex molecular processes. This commentary provides a discussion and critical analysis of the core AIMMD framework, explores its recent extensions, and offers an assessment of the method's potential impact and limitations.

Background

Elucidating complex reaction mechanisms remains a central challenge in both chemical and biological sciences. Molecular dynamics (MD) simulations have long served as a powerful tool for uncovering mechanistic details, as they offer atomistic resolution and direct access to dynamical information. However, many reactions of interest are rare events that occur on timescales far longer than those accessible to straightforward MD, making direct sampling of reactive transitions prohibitively inefficient. To overcome this limitation, numerous enhanced sampling techniques have been developed, including biased approaches (e.g., umbrella sampling [2], metadynamics [3, 4]) and unbiased approaches (e.g., path sampling [5–10], weighted ensemble [11]). Critically, the efficiency of these methods depends strongly on their adherence to the reaction coordinate (RC), which is unknown *a priori*. The RC is a special coordinate of the system that accurately captures the dynamical progress from the reactant (state A) to the product (state B) state. [12] The committor probability ($p_B(x)$), defined as the probability that a trajectory initiated from configuration x reaches state B before returning to state A , is widely regarded as the optimal RC [12, 13]. Yet, computing $p_B(x)$ directly is computationally demanding, as it requires shooting hundreds of short trajectories from numerous configurations. Furthermore, the p_B does not directly provide any physical insights into the structural and dynamical properties important for the transition.

Transition path sampling (TPS) provides an unbiased framework for generating reactive trajectories without requiring prior knowledge of the RC. In short, TPS performs Monte Carlo (MC) moves in path space, iteratively modifying a configuration of an existing path to generate a new path. In a typical move, a configuration (also

known as the shooting point) along the current reactive trajectory is selected, its momenta are perturbed, and the forward and backward integrations generate a proposed new path [5, 14]. If the new trajectory connects state A to state B, it is accepted; otherwise, it is rejected. The algorithm can be summarized in Algorithm 1:

Algorithm 1 Logic of the TPS Algorithm

Require: One reactive trajectory \mathcal{TP} that connects state A to state B

```

1: while sampling is not converged do
2:   Select a shooting point  $\mathbf{x}_s$  from the current  $\mathcal{TP}$  according to a probability  $P_{\text{sel}}(\mathbf{x}_s | \mathcal{TP})$ 
3:   Perturb the momenta at  $\mathbf{x}_s$  to  $\mathbf{x}_s^n$  according to  $P_{\text{acc}}(\mathbf{x}_s \rightarrow \mathbf{x}_s^n)$ 
4:   if not reject then
5:     while not in B do
6:       Integrate equations of motion forward in time from  $\mathbf{x}_s^n$ 
7:       if in A before in B then
8:         reject = True
9:       Retain the previous trajectory  $\mathcal{TP}$ 
10:      Go back to Step 2
11:    end if
12:   end while
13:   while not in A do
14:     Integrate equations of motion backward in time from  $\mathbf{x}_s^n$ 
15:     if in B before in A then
16:       reject = True
17:     Retain the previous trajectory  $\mathcal{TP}$ 
18:     Go back to Step 2
19:   end if
20:   end while
21:   Combine the forward and backward segments and record the new trajectory  $\mathcal{TP}'$ 
22:   Update  $\mathcal{TP} \leftarrow \mathcal{TP}'$ 
23: end if
24: end while
  
```

Although TPS bypasses the need for the RC when sampling, its efficiency still depends on choosing a reasonable collective variable (CV)—a physical quantity of the system—that can distinguish between states A and B. TPS can suffer from poor acceptance rates or inadequate exploration when the chosen CV does not correlate well with the progress of the reaction. Thus, identifying an effective CV or the RC remains a major bottleneck.

Artificial Intelligence for Molecular Mechanism Discovery (AIMMD) addresses this bottleneck by integrating machine learning (ML) with the TPS framework. Rather than relying on a fixed CV, AIMMD learns the RC as the committer on-the-fly while performing the ML-enhanced TPS. This on-the-fly learning and adaptation of the RC effectively mitigates the issues associated with a poorly chosen sampling coordinate, allowing the algorithm to focus computational resources on the transition region, which is critical for generating transition paths (TP) in TPS. Furthermore, AIMMD elevates the scientific utility of the committer by employing symbolic regression to describe it as an analytical function of preselected CVs, thereby expanding its interpretability. Overall, AIMMD showcases a robust framework for the efficient discovery and characterization of complex molecular pathways.

Overview of AIMMD

The AIMMD algorithm is initiated in a manner analogous to TPS. First, states A and B of the system are defined along a physically motivated CV. An initial TP connecting A and B is then generated to jump-start

TPS. After performing several TPS moves, the shooting points and their outcomes from this initial run provide the training data for the first ML cycle.




Figure 1: **Overview of AIMMD Schematic.** First, AIMMD is trained in a self-consistent manner, where the process alternates between rounds of sampling using TPS and training the neural network. Once trained, the resulting committor function is described in terms of the physical CVs of the system through symbolic regression. Figure reproduced with permission from Figure 1a of Ref. 1.

AIMMD employs a feed-forward neural network in which each data point \mathbf{x} corresponds to a molecular configuration (i.e., a shooting point) described by N physical CVs (i.e., features). For example, in an ion association–dissociation process, such CVs may include the inter-ionic distance or the coordination numbers of surrounding solvent atoms. These inputs are nonlinearly combined as they propagate through the network to output the RC in the form of a logit-committor, $q(\mathbf{x}|\boldsymbol{\theta})$. The committor is then given by Eq. 1:

$$p_B(\mathbf{x}) = \frac{1}{1 + e^{-q(\mathbf{x}|\boldsymbol{\theta})}} \quad (1)$$

The network parameters $\boldsymbol{\theta}$ are learned from the successes and failures of the most recent k shooting attempts. Training is performed by minimizing the negative log-likelihood loss,

$$L = - \sum_{i=1}^k \left[n_i^A \ln(1 - p_B(\mathbf{x}_i)) + n_i^B \ln p_B(\mathbf{x}_i) \right] = - \sum_{i=1}^k \left[n_i^A \ln\left(1 - \frac{1}{1 + e^{-q(\mathbf{x}_i|\boldsymbol{\theta})}}\right) - n_i^B \ln\left(1 + e^{-q(\mathbf{x}_i|\boldsymbol{\theta})}\right) \right] \quad (2)$$

where n_i^A and n_i^B encode the outcome of shooting point \mathbf{x}_i . For a reactive path (AB), $n_i^A = n_i^B = 1$. For nonreactive paths, $n_i^A = 2, n_i^B = 0$ for AA trajectories and $n_i^A = 0, n_i^B = 2$ for BB trajectories. Minimizing this loss, therefore, yields network parameters that maximize the likelihood of the observed shooting outcomes.

AIMMD proceeds in a self-consistent loop, where TPS sampling and ML training alternate (Fig. 1). After each round of model training, TPS continues by selecting a shooting point on the most recently generated TP according to a Lorentzian distribution centered on the transition-state ensemble (TSE):

$$P_{\text{sel}}(\mathbf{x}|\mathcal{TP}) = \frac{1}{\sum_{\mathbf{x}' \in \mathcal{TP}} \left[\frac{q(\mathbf{x})^2 + \gamma^2}{q(\mathbf{x}')^2 + \gamma^2} \right]}, \quad (3)$$

where larger values of γ promote broader exploration. This scheme enhances sampling efficiency by favoring shooting points near the transition state, while still permitting occasional excursions away from it.

The training cycle terminates once the committor model is deemed sufficient, as determined by whether the expected number of TPs from the last k shooting points matches the number actually observed. The expected

number of TPs is given by

$$n_{TP}^{\text{exp}} = \sum_{i=1}^k 2(1 - p_B(\mathbf{x}_i))p_B(\mathbf{x}_i) \quad (4)$$

where $p_B(\mathbf{x}_i)$ is the predicted probability of reaching state B before the shooting outcome at step i is known. Once trained, the committor can be expressed in an interpretable form using symbolic regression, which represents the learned RC as a combination of known physical CVs. Although the degree of interpretability depends on how complex the resulting functional form is, this representation still provides insights into the relative importance and interactions of the underlying CVs of the system.

In Jung *et al.*[1], AIMMD was applied to a range of systems, including ion association/dissociation in solution, gas hydrate nucleation, homopolymer folding, and the assembly of the transmembrane protein Mga2 in a lipid bilayer. Notably, in the case of ion association/dissociation, Jung *et al.* [1] demonstrated transferability of the learned mechanism. Starting from a model trained for LiCl, they performed transfer learning by retraining only the final neural network layer to study the ion association-dissociation of other monovalent salts. In the Mga2 assembly study, they further showcased the use of parallel TPS simulations to generate pooled data for model training. This strategy enabled the identification of two distinct reaction pathways and produced a committor landscape that reflects their coexistence.

Commentary and Critical Analysis

While AIMMD presents a promising approach for automated RC discovery and efficient rare event sampling, our examination highlights several areas where methodological rigor and further development would strengthen the framework. The method clearly demonstrates its ability to identify an RC that correlates with the committor, thereby facilitating the localization of the TSE and improving the efficiency of generating TPs within TPS. However, a notable gap remains regarding a rigorous assessment of the quality of the learned committor. Although Jung *et al.* showed that the predicted committor aligns well with numerically estimated committor values, the framework itself does not incorporate a definitive procedure for verifying that the learned RC is the true committor. A well-established validation tool—the histogram test [12, 15]—provides such an assessment by checking whether configurations predicted to have the same committor indeed exhibit statistically identical outcomes in short shooting trials. Integrating this test into AIMMD would provide a quantitative evaluation of whether the learned function accurately represents the true committor.

A second issue concerns the sampling imbalance inherent in the learning procedure. Although AIMMD employs a Lorentzian distribution for P_{sel} to balance exploration and exploitation, the majority of shooting points chosen would be configurations near the transition state, which generates TPs. Consequently, the training data are disproportionately concentrated around configurations with $p_B(\mathbf{x}) \approx 0.5$. This uneven sampling reduces the predictive accuracy of the learned committor for regions farther from the TSE (i.e., for configurations with $p_B < 0.5$ or $p_B > 0.5$), where fewer data points are available. In principle, this could hinder downstream tasks that rely on accurate committor predictions away from the TSE. Therefore, highlighting strategies to adaptively broaden the sampling distribution or to include additional exploratory shooting moves would enhance the robustness of the approach.

Furthermore, AIMMD inherits a key limitation of TPS itself, which is the difficulty in thoroughly exploring multiple reaction pathways when the system exhibits branching mechanisms. Accessing distinct pathways requires generating a shooting point near the dividing region separating the pathways, which may be rare if the energy barrier between pathways is non-negligible. As a result, the probability of discovering new pathways depends sensitively on both the barrier height and the degree of overlap between the respective TSEs. In contrast, methods such as replica exchange transition interface sampling (RETIS) [7] employ exchanges between neighboring interface ensembles, allowing trajectories to cross barriers in "pathway space" more readily. AIMMD partially addresses this limitation through training on data from multiple, parallel TPS instances, as demonstrated in the Mga2 study. However, the degree to which this strategy mitigates sampling bias depends strongly on how those parallel trajectories are initialized. A more detailed discussion of how AIMMD performs

when confronted with reaction networks containing multiple, well-separated pathways would benefit future users aiming to apply the method in such settings.

Extensions of AIMMD

The AIMMD framework has already motivated several methodological extensions that broaden its utility and scope. An early extension of AIMMD is the waste-recycling TPS scheme introduced by Lazzeri *et al.* [16]. In standard TPS, only successful TPs are used for analysis, whereas nonreactive (AA and BB) trajectories are discarded. The waste-recycling variant utilizes these rejected paths, together with data from additional equilibrium simulations, to extract information from a more extensive sampling of the configuration space. Incorporating this broader dataset yields a more complete characterization of the reaction landscape, enabling improved estimation of both the rate constant and the free energy surface as a function of the learned committor. Furthermore, in this AIMMD-TPS implementation, shooting points are chosen to balance the *exploration* of the relevant configurational region with *exploitation* of points likely to generate reactive trajectories by using a uniform distribution as a function of the committor instead of the Lorentzian. This strategy helps alleviate the imbalance in training data that arises when most samples originate from near the transition region, thereby improving the committor estimates for configurations farther from the transition state.

A more recent extension integrates AIMMD with Transition Interface Sampling (AIMMD-TIS) [17]. In this approach, an initial AIMMD-TPS run provides a preliminary estimate of the logit-committor $q(x)$, which is then used to determine approximate boundaries for states A and B and to place TIS interfaces that serve as analogs of isocommittor surfaces. Since TIS saves both reactive and nonreactive trajectories, it generates a much more diverse ensemble of configurations than TPS alone. This expanded sampling substantially improves the coverage of the committor across configuration space, particularly in regions away from the TSE where AIMMD-TPS may lack sufficient data. At convergence, AIMMD-TIS yields not only the reaction rate and free energy surface, but also a refined committor model and an associated set of isocommittor surfaces. The latter enables the computation of the mean gradient of the logit-committor with respect to each neural network input feature, thereby providing a systematic and quantitative measure of feature importance at different stages along the reaction pathway.

Furthermore, by leveraging the interface structure of TIS, AIMMD-TIS enhances the exploration of alternative reaction pathways compared to AIMMD-TPS. The use of multiple interfaces, each generating their own ensemble of trajectories, increases the likelihood of sampling distinct mechanistic channels. While this makes AIMMD-TIS a better choice for systems with multiple pathways or mechanistic heterogeneity compared to AIMMD-TPS, an approach incorporating RETIS with AIMMD would provide an even more comprehensive exploration of pathway diversity.

Conclusion

Historically, the broader adoption of path sampling methods has lagged behind that of biased techniques such as metadynamics, in part due to limited software accessibility. This gap is now narrowing with the development of user-friendly tools such as OpenPathSampling (OPS) [18, 19], AIMMD's own software package [1], and PyRETIS [20–22], which collectively lower the barrier to entry for practitioners.

A further practical challenge arises from the inherently iterative nature of TPS-based approaches. For highly diffusive processes, such as salt nucleation, reactive trajectories can be extremely long, which raises the question about the efficiency of repeatedly generating such trajectories during RC refinement. The extent to which AIMMD can maintain its efficiency in these slow dynamical regimes remains to be fully tested. Nevertheless, if the RC can be refined significantly using only a modest number of trajectories, as demonstrated in the ion association–dissociation case through few-shot learning, AIMMD has the potential to remain highly effective even in complex, slow dynamical systems.

Overall, AIMMD serves as a valuable tool for computational chemistry, providing an automated and interpretable framework for sampling and elucidating complex pathways. By combining ML with TPS, AIMMD

not only enhances sampling efficiency but also yields physically meaningful, interpretable RC through symbolic regression. Through directly addressing one of the most pressing computational bottlenecks in simulations, it has the potential to rapidly enhance the understanding of complex systems.

Acknowledgments

This work was supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Award No. DE-SC0025496. P.M. acknowledges support from Wayland E. Noland via the Wayland E. Noland Excellence Fellowship in Chemistry.

References

- [1] Jung H, Covino R, Arjun A, et al. "Machine-guided path sampling to discover mechanisms of molecular self-organization". *Nature Computational Science*. 2023, 3:334–345.
- [2] Torrie G, Valleau J. "Nonphysical sampling distributions in monte carlo free-energy estimation: Umbrella sampling". *Journal of Computational Physics*. 1977, 23(2):187–199.
- [3] Laio A, Parrinello M. "Escaping free-energy minima". *Proceedings of the National Academy of Sciences*. 2002, 99(20):12562–12566.
- [4] Barducci A, Bussi G, Parrinello M. "Well-tempered metadynamics: A smoothly converging and tunable free-energy method". *Phys Rev Lett*. 2008, 100:020603.
- [5] Bolhuis PG, Chandler D, Dellago C, et al. "Transition path sampling: Throwing ropes over rough mountain passes, in the dark". *Annual Review of Physical Chemistry*. 2002, 53:291–318.
- [6] van Erp TS, Bolhuis PG. "Elaborating transition interface sampling methods". *Journal of Computational Physics*. 2005, 205(1):157–181.
- [7] Cabriolu R, Skjelbred Refsnes KM, Bolhuis PG, et al. "Foundations and latest advances in replica exchange transition interface sampling". *The Journal of Chemical Physics*. 2017, 147(15):152722.
- [8] Allen RJ, Warren PB, ten Wolde PR. "Sampling rare switching events in biochemical networks". *Phys Rev Lett*. 2005, 94:018104.
- [9] DeFever RS, Sarupria S. "Contour forward flux sampling: Sampling rare events along multiple collective variables". *J Chem Phys*. 2019, 150(2):024103.
- [10] Hall SW, Díaz Leines G, Sarupria S, et al. "Practical guide to replica exchange transition interface sampling and forward flux sampling". *J Chem Phys*. 2022, 156(20):200901.
- [11] Zwier MC, Adelman JL, Kaus JW, et al. "Westpa: An interoperable, highly scalable software package for weighted ensemble simulation and analysis". *Journal of Chemical Theory and Computation*. 2015, 11(2):800–809. PMID: 26392815.
- [12] Peters B. "Reaction Coordinates and Mechanistic Hypothesis Tests". *Annu Rev Phys Chem*. 2016, 67:669–690.
- [13] Liang Y, Díaz Leines G, Drautz R, et al. "Identification of a multi-dimensional reaction coordinate for crystal nucleation in Ni3Al". *The Journal of Chemical Physics*. 2020, 152(22):224504.
- [14] Bolhuis PG, Dellago C. Trajectory-based rare event simulations. John Wiley & Sons, Ltd, 2010. Chapter 3, p. 111–210. <https://onlinelibrary.wiley.com/doi/10.1002/9780470890905.ch3>.
- [15] Peters B. "Using the histogram test to quantify reaction coordinate error". *The Journal of Chemical Physics*. 2006, 125(24):241101.
- [16] Lazzeri G, Jung H, Bolhuis PG, et al. "Molecular free energies, rates, and mechanisms from data-efficient path sampling simulations". *Journal of Chemical Theory and Computation*. 2023, 19(24):9060–9076. PMID: 37988412.

- [17] Breebaart RS, Lazzeri G, Covino R, et al. "Understanding reaction mechanisms from start to finish", 2025. <https://arxiv.org/abs/2507.04052>.
- [18] Swenson DWH, Prinz JH, Noe F, et al. "OpenPathSampling: A Python framework for path sampling simulations. 1. Basics". *Journal of Chemical Theory and Computation*. 2019, 15(2):813–836. PMID: 30336030.
- [19] Swenson DWH, Prinz JH, Noe F, et al. "OpenPathSampling: A Python framework for path sampling simulations. 2. Building and customizing path ensembles and sample schemes". *Journal of Chemical Theory and Computation*. 2019, 15(2):837–856.
- [20] Lervik A, Riccardi E, van Erp TS. "Pyretis: A well-done, medium-sized python library for rare events". *Journal of Computational Chemistry*. 2017, 38(28):2439–2451.
- [21] Riccardi E, Lervik A, Roet S, et al. "Pyretis 2: An improbability drive for rare events". *Journal of Computational Chemistry*. 2020, 41(4):370–377.
- [22] Vervust W, Zhang DT, Ghysels A, et al. "Pyretis 3: Conquering rare and slow events without boundaries". *Journal of Computational Chemistry*. 2024, 45(15):1224–1234.