Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

All Paths All at Once: A Quantum Monte Carlo and Neural Committor Approach to Reaction Pathway Search

Chien-Chai Chang

frank20040506@gapp.nthu.edu.tw

National Tsing Hua University https://orcid.org/0009-0001-5314-7976

Research Article

Keywords: Quantum Computing, Physics Computation, Reaction Pathway Search

Posted Date: November 2nd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11557561/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

All Paths All at Once: A Quantum Monte Carlo and Neural Committor Approach to Reaction Pathway Search

Chien-Chai Chang
Department of Engineering and System Science,
National Tsing Hua University (NTHU), Hsinchu City, 300044, Taiwan

Locating competing reaction pathways is difficult when potential-energy gradients are unavailable. Here we present tReax_expedition, an energy-only approach combining coupled Monte Carlo exploration, stochastic quench labelling and restrained continuation. Exploration identifies connected states, while continuation resolves candidate channels and records sampled chain maxima. Single-search benchmarks on three two-dimensional potentials recovered an intermediate and competing routes without transition-state guesses. For eight of nine gas-phase reactions on GFN2-xTB, sampled chain maxima differed from independently refined saddle energies by less than 0.76 kJ mol^{-1}. Alanine dipeptide instead revealed incomplete conformational relaxation, showing that route connectivity can persist despite structural and energetic errors. Neural committor learning and local straddling-pair sampling connect the resulting paths to transition-region sampling, although molecular shooting diagnostics indicate limited committor calibration. These results show how energy-only exploration and continuation recover candidate pathways, while distinguishing route connectivity from geometric localization, barrier resolution and finite-temperature transition statistics.

1. INTRODUCTION

Transition states determine the rates and mechanisms of chemical reactions and conformational changes[1–3]. Locating them requires searching rarely visited regions of a potential energy surface, often along several competing channels[4, 5]. This task becomes expensive when accurate electronic-structure methods provide energies without analytic gradients[6].

Most transition-state searches use forces. Single-ended methods locate a saddle from one structure[7, 8], whereas chain-of-states methods relax a path between prescribed endpoints[9–12]. Their outcomes depend on the starting structure or path. Generative models instead propose transition states from endpoint structures[13, 14], molecular graphs[15] or metastable configurations under guided diffusion[16]. These proposals depend on training data and require validation on the target energy surface.

Rare-event methods describe the transition through the committor, the probability of reaching a product basin before a reactant basin[17–19]. Its isosurface $q = 1/2$ defines a transition-state ensemble even when several channels compete[4]. Neural committors can be learned from path-sampling outcomes[20, 21] or variational formulations[22–25]. Kang et al. used a committor-derived Kolmogorov bias to sample the transition region while learning the committor[26]. Trizio et al. combined this bias with on-the-fly probability enhanced sampling to sample basins, intermediates and competing transition channels[27]. Their molecular-dynamics implementation requires forces and specified endpoint basins. Learned slow modes[28], descriptor-free graph collective variables[29] and graph committors[25] reduce reliance on manually chosen inputs. These advances do not by themselves define the energy-only exploration and state-discovery procedure considered here.

An energy-only alternative is the Monte Carlo transition-state search (MCTSS) of Tikhonov and Santra[6]. Inspired by real-space diffusion Monte Carlo[30], it drives walkers from opposite endpoints until they meet. The authors demonstrated reaction-path searches with DLPNO-CCSD(T) energies[6]. MCTSS was designed for search rather than equilibrium sampling: its acceptance includes a kinetic factor, and its geometric pre-selection redraws trials until one passes. Its time parameter also uses harmonic frequencies, while its barriers are estimated from binned trajectory averages. These choices motivate a statistically explicit energy-only search.

Here we present tReax_expedition as a framework for energy-only candidate-route search. Coupled-walker exploration supplies paths for quench-based state assignment and restrained continuation chains. The chains resolve candidate channels and record their largest evaluated energies; the continuous-path bound retains an explicit discretization correction. Committor learning then combines a variational loss with shooting trajectories generated by the same energy-only kernel. Supplementary Table S.1 specifies which statements are exact, which require modelling assumptions and which depend on numerical heuristics.

For isolated molecules with invariant energies and biases, tReax_expedition respects $SE(3)$ symmetry. Before frame removal, its molecular sampling kernel is equivariant in distribution under proper rotations and translations. The scalar committor is invariant, and its Cartesian gradient is rotation-equivariant. Frame removal preserves the process of invariant observables; coupled walkers retain a common frame (Section 4.1; Supple-

mentary Section 2). This claim excludes fixed-frame periodic systems, constrained atoms and the two-dimensional model potentials. Supplementary Table S.5 defines the benchmark systems; Supplementary Table S.6 gives their descriptors and shooting budgets.

2. RESULTS

2.1. Benchmark systems and reference quantities

We evaluated channel recovery, chain energies and transition-region structures separately from finite-temperature committor prediction. The benchmark comprised three two-dimensional potentials, nine gas-phase reactions and alanine dipeptide (Supplementary Table S.5). The two-dimensional models tested a single channel, an intermediate and competing routes[6, 27, 31]. The gas-phase set combined eight Baker–Chan reactions in the SACTS input geometries[32, 33] with $Cl^- + CH_3Br \rightarrow CH_3Cl + Br^-$. All nine gas-phase searches and their reference comparisons used GFN2-xTB[34]. Alanine dipeptide tested the $C7_{eq} \rightarrow C7_{ax}$ conformational transition on the energy-only ff99SB surface[35]. The searches used endpoint structures, without a transition-state guess or potential-energy gradients.

Two candidate structures were assessed for each retained channel: its highest chain node and the crossing located by bisecting quench labels. The latter, $\mathbf{R}_0^{\ddagger}$, estimates a boundary between stochastic quench basins; it does not establish the finite-temperature isocommittor $q = 1/2$. We compared both candidates with analytic-model saddles or independently refined molecular saddles on the same energy surface. Molecular references were checked for one imaginary mode; their refinement and Hessian calculations were used only for evaluation. Hessian analysis characterizes the local saddle but does not establish its connection to both intended endpoints. Structural deviations were measured after alignment and symmetry handling (Table I). Saddle-referenced deviations quantify localization error, whereas Table I reports activation estimates relative to a common reconstructed state-A energy.

The chain quantity E_{up} includes evaluated link, quench-trajectory and endpoint-descent energies (Methods Section 4.3). Certification checks the endpoint labels and these evaluated energies. A continuous-path upper bound additionally requires the uncomputed discretization correction in Supplementary Section 4.2. The final structural benchmark contained one run per system, so its cross-system counts do not estimate success probabilities or seed variability.

2.2. Energy-only chains recovered intermediates and competing routes

The two-dimensional tests recovered the expected connectivity without supplying an intermediate or a transition-state structure. On the proton-transfer surface, one chain connected the two wells through the saddle region (Fig. 1a). On the Müller–Brown surface, quench labelling identified the intermediate and separated the route into $A \rightarrow I_1$ and $I_1 \rightarrow B$ chains (Fig. 1b). The two-path surface yielded both its lower and upper channels (Fig. 1c). These observations establish recovery of the distinct routes present in this run, without establishing exhaustive channel discovery.

Chain-energy agreement and crossing geometry gave different measures of resolution. Across the five two-dimensional saddle comparisons, $E_{up} - E_{saddle}$ ranged from -0.002 to $+0.025$ kJ mol^{-1} (Fig. 1). Crossing displacements ranged from 0.004 to 0.101 Å. The largest displacement occurred on the nearly flat lower channel of the two-path surface, despite a crossing-energy difference of only 0.006 kJ mol^{-1}. Energy agreement therefore did not imply equally precise geometric localization. Small negative differences on the proton-transfer and Müller–Brown surfaces also exposed the finite resolution of the evaluated chains and quench-label bisection.

2.3. Most gas-phase chains approached the reference saddle energies

Eight of the nine GFN2-xTB reactions had $|E_{up} - E_{saddle}| < 0.76$ kJ mol^{-1} (Table I). Their crossing RMSDs were below 0.019 Å. The profiles place the crossings near the principal chain maximum across hydrogen transfer, dissociation, addition and ring opening (Fig. 2a–i). The orange markers show the highest node, whose energy need not set E_{up}. Selected interatomic distances in the reference and crossing structures provide a local view of this geometric agreement (Supplementary Fig. S.2).

Acetaldehyde tautomerization was the exception. Its chain estimate exceeded the reference saddle energy by 6.08 kJ mol^{-1}, whereas the crossing lay 3.48 kJ mol^{-1} above it (Fig. 2f). Bisection reduced the RMSD from 0.164 Å at the highest node to 0.043 Å at the crossing. Similar improvements occurred for S_N2 substitution (0.040 to 0.018 Å) and HF addition to ethylene (0.023 to 0.013 Å). Thus, refining the label-change interval improved localization in these cases, but did not remove the residual energy error for acetaldehyde. Its additional retained chain was symmetry-equivalent and did not establish a different mechanism.

Endpoint relaxation gave a separate energetic check: the nine gas-phase reaction energies differed from independently refined minima by less than 0.08 kJ mol^{-1} (Table I). This agreement concerns endpoint energies and does not establish saddle accuracy. These comparisons measure search accuracy on the specified semiempirical surface, rather than accuracy against experimental barriers or higher-level electronic structure. The HF-3c endpoint geometries supplied for S_N2 were not GFN2-xTB minima, so published HF-3c barriers were not used to score this benchmark. Likewise, the slightly negative HCN chain difference (-0.013 kJ mol^{-1}) does not con-

stitute a verified lower-energy mechanism. It is consistent with a discretely sampled chain undershooting the reference saddle. Across systems, the joint comparison retained structural displacement and energy deviation rather than treating either as a proxy for the other (Supplementary Fig. S.1).

2.4. Alanine dipeptide exposed incomplete conformational relaxation

The structural and energetic agreement of the small-molecule tests did not extend to alanine dipeptide. Four chains connected the prescribed conformational states, but their E_{up} values exceeded the refined ff99SB saddle by 11.57–25.00 kJ mol^{-1}. On the lowest chain, the crossing had all-atom and heavy-atom RMSDs of 0.453 and 0.284 Å, respectively (Fig. 3b). Its energy remained 5.13 kJ mol^{-1} above the reference (Fig. 3c).

The dihedral projection localized the unresolved conformational direction. The lowest-chain crossing had $(\phi, \psi) = (-2.8°, 5.4°)$, compared with $(-2.1°, -26.4°)$ at the reference saddle. Across all four crossings, ψ remained between 5.4° and 12.6° (Fig. 3a). This systematic displacement is consistent with incomplete relaxation along the soft conformational direction, although these data do not isolate its cause. Comparison with the unrelaxed transition-state guess gave large negative energy differences and would have concealed this error. The refined same-surface saddle therefore supplied the relevant evaluation reference. Endpoint energies exposed a related limitation: ΔE was 0.122 eV, compared with 0.062 eV for the refined minima (Table I), an excess of 5.78 kJ mol^{-1} associated with greater residual energy in state B.

2.5. Committor validation remained a separate requirement

Recovering a saddle neighbourhood did not by itself validate a learned committor. The grid reference describes overdamped diffusion with absorbing discs, whereas shooting follows the finite-step Monte Carlo kernel and its prescribed cores. Their agreement is therefore a diagnostic, rather than an exact validation of the Monte Carlo hitting probability.

Molecular shooting tests exposed a narrower validation range. In a development test using the shooting likelihood, the HCOH crossing prediction was 0.43, against a measured product fraction of 0.60 ± 0.11. HCN predictions remained poorly calibrated. For alanine dipeptide, 14–20 of 20 probe trajectories failed to reach either core within the prescribed budget. Excluding these trajectories prevents an unconditional committor estimate without assumptions about their eventual outcomes. The molecular results therefore support quench-label localization more strongly than finite-temperature committor calibration. Supplementary Fig. S.3 distinguishes the broad biased sampling used for learning from the local midpoint ensemble around a crossing. Its latent projections illustrate their different support; they do not extend the molecular calibration evidence.

This capability also required substantial energy evaluation. The nine gas-phase searches used 0.249–1.517 million oracle calls per system; alanine dipeptide used 3.254 million, for the specified search budgets.

3. DISCUSSION

The present results support separating energy-only pathway search from the statistical characterization of a transition. Coupled-walker exploration proposes routes, quench labels identify states, and continuation tests their connectivity and sampled energy maxima. Committor learning and local ensemble sampling then address different questions. This division retains the exploratory role of MCTSS[6] while connecting its output to variational committor learning and enhanced sampling[20, 26, 27]. The searches recover candidate channels without potential-energy gradients. Channel discovery remains limited by exploration and the geometric grouping rule.

Structural localization, chain energy and endpoint relaxation provide complementary checks. A nearly flat saddle region can tolerate appreciable displacement with little energy change, whereas a good endpoint reaction energy says little about the intervening barrier. The model and molecular comparisons expose both distinctions (Figs. 1–3; Table I). Alanine further shows that an identifiable route can coexist with incomplete relaxation of both the crossing and the endpoint states. Its errors are consistent with limited exploration of soft conformational directions, but the present runs do not isolate whether proposals, coordinates or quench budgets dominate. These observations favor evaluating connectivity, geometry and energetics separately rather than accepting a single agreement metric as evidence of convergence.

The distinction between broad sampling and the local straddling-pair ensemble (labelled TSE in code and figures) is equally consequential. The biased walkers used for committor learning combine a Kolmogorov bias with a bias along the network latent, following the complementary roles of transition-region enrichment and basin-to-basin exploration in earlier work[26, 27]. Their broad latent coverage supplies configurations for reweighted equilibrium projections (Supplementary Fig. S.3c,d). Known frozen biases make multistate Bennett acceptance ratio (MBAR) reweighting possible[36]; they do not ensure equilibration, overlap between sampled ensembles or useful effective sample size. Moreover, the percentile cap used during training alters the estimator. The displayed projections therefore illustrate the sampling objective without establishing an uncertainty-controlled free-energy calculation.

The straddling-pair ensemble instead concentrates

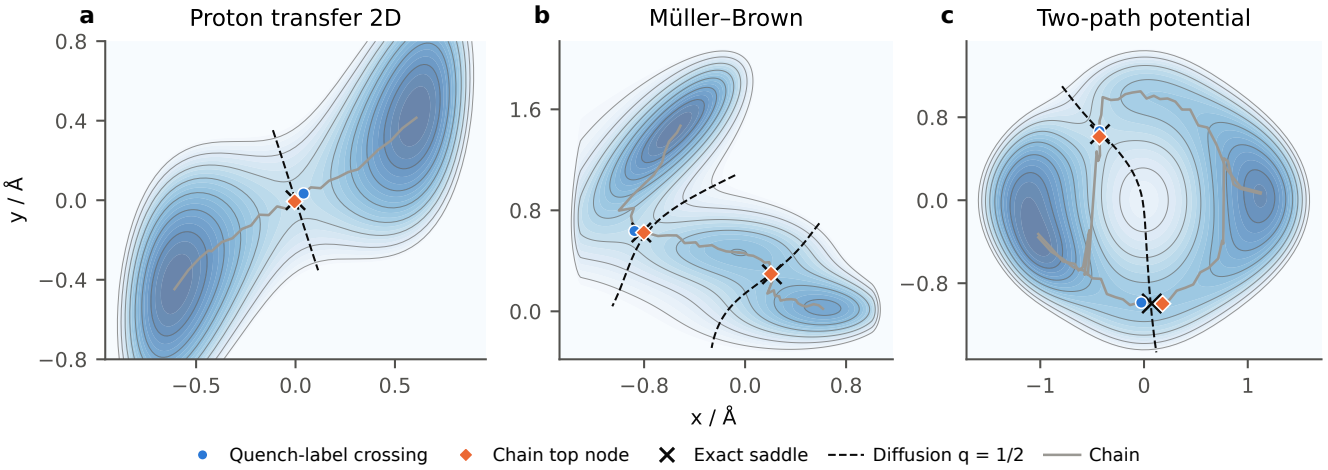


FIG. 1. **Recovery of an intermediate and competing reaction channels.** **a**, Proton transfer coupled to a harmonic bath. **b**, Müller–Brown potential. **c**, Two-path potential. Grey curves trace retained continuation nodes over potential-energy contours; blue circles denote quench-label crossings, orange diamonds the highest nodes and black crosses the reference saddles. Dashed lines are the $q = 1/2$ contours of numerical diffusion references, not learned committors. Coordinates are in Å. Each model was searched once; multiple channels within a search are not independent replicates.

TABLE I. **Same-surface activation energies, reaction energies and candidate-structure comparisons.** Both E_a columns use the recorded state-A energy: $E_{\mathrm{saddle}} - E_A$ for the independently refined first-order saddle and $E_{\mathrm{up}} - E_A$ for tReax_expedition. Reference ΔE is the product-minus-reactant energy of independently refined minima; tReax_expedition $\Delta E = E_B - E_A$ uses the final recorded state minima. All energies are in eV, without zero-point or thermal corrections; ΔE is a potential reaction energy, not a finite-temperature enthalpy. Crossing RMSD compares the quench-label crossing with the local reference saddle after alignment and symmetry handling; this comparison measures geometric agreement with a local saddle reference. Gas-phase reactions use GFN2-xTB; alanine dipeptide uses ff99SB. Values are rounded to three decimals, with one search per system. Endpoint-energy sources are listed in Data availability.

		Reference		tReax_expedition		
Reaction	No. of atoms	E_a (eV)	ΔE (eV)	E_a (eV)	ΔE (eV)	Crossing RMSD (Å)
$HCN \rightarrow HNC$	3	3.175	0.868	3.175	0.868	0.013
$H_2CO \rightarrow H_2 + CO$	4	3.167	1.862	3.168	1.863	0.003
$HCOH \rightarrow H_2CO$	4	1.744	-2.244	1.745	-2.244	0.008
$HCCH \rightarrow H_2CC$	4	2.614	2.360	2.614	2.360	0.008
$CH_3O^{\bullet} \rightarrow {}^{\bullet}CH_2OH$	5	1.723	0.370	1.726	0.370	0.012
$Cl^- + CH_3Br \rightarrow CH_3Cl + Br^-$	6	0.275	-0.117	0.283	-0.116	0.018
$CH_3CHO \rightarrow CH_2CHOH$	7	2.918	0.248	2.981	0.249	0.043
$C_2H_4 + HF \rightarrow CH_3CH_2F$	8	1.523	-1.606	1.528	-1.606	0.013
$C_3H_5O_2^{\bullet}$ ring opening	10	0.792	0.528	0.800	0.528	0.016
Alanine dipeptide, $C7_{eq} \rightarrow C7_{ax}$	22	0.347	0.062	0.467	0.122	0.453

midpoint configurations near a surface fixed by a quench-label crossing (Supplementary Fig. S.3b,c). In these illustrated runs, the side constraint uses the chain coordinate rather than the learned committor. Midpoint reweighting corrects the pair-energy factor, while the side constraint retains a finite-width window and, in the stiff limit, a coordinate-dependent surface measure (Supplementary Section 5.9). This ensemble can characterize the sampled neighbourhood without determining basin populations or a complete free-energy profile along the latent. The blue intervals in Supplementary Fig. S.3d indicate its restricted support, not independent barrier estimates. Proximity to a saddle also does not establish that this surface is the finite-temperature Monte Carlo isocommittor or a distribution of reactive flux.

The statistical interpretation remains tied to the specified sampling dynamics. The network targets a first-hitting probability of the finite-step Monte Carlo kernel; the variational functional and the plotted grid references describe a diffusion approximation. Agreement with those references (Supplementary Fig. S.3a) is informative but does not replace shooting under the actual

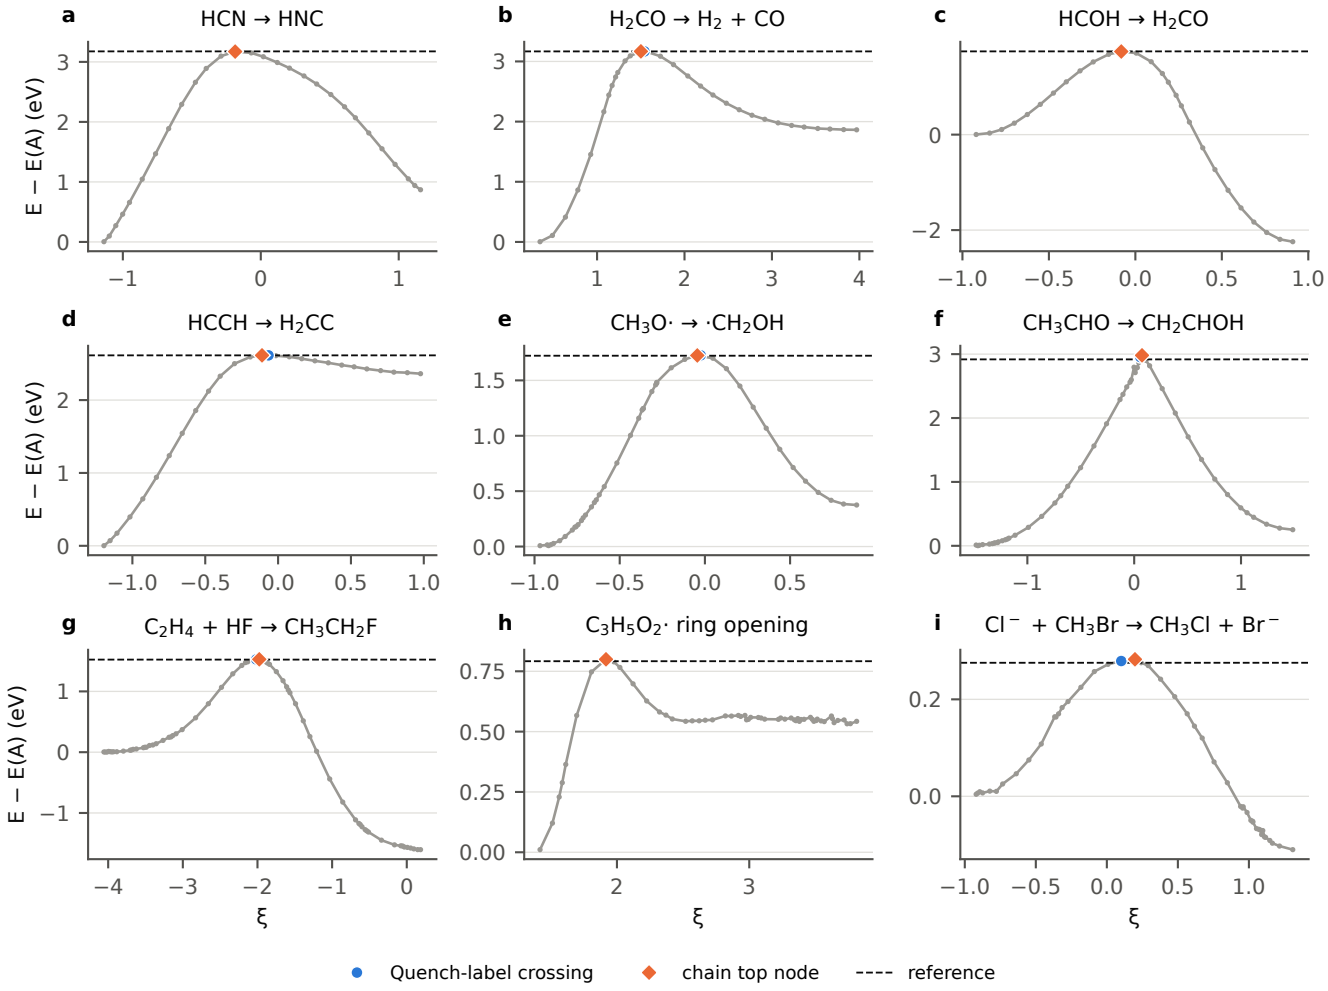


FIG. 2. **Gas-phase reaction profiles reconstructed from energy-only continuation. a–i**, Primary continuation chains for HCN isomerization, formaldehyde dissociation, HCOH isomerization, acetylene rearrangement, methoxy isomerization, acetaldehyde tautomerization, HF addition to ethylene, radical ring opening and S_N2 substitution, respectively. Grey points show node energies relative to the lowest recorded state-A minimum, plotted against ξ in Å, the summed breaking-bond lengths minus the summed forming-bond lengths (Supplementary Section 5.6). Blue circles mark quench-label crossings, orange diamonds the highest nodes and dashed lines the independently refined saddle energies. All energies were evaluated with GFN2-xTB. These potential-energy profiles represent one search per reaction; they are neither free-energy profiles nor force-verified minimum-energy paths. Node maxima need not equal E_{up}, which also includes link and quench-trajectory evaluations.

kernel and core definitions. Non-committing molecular shots further restrict what can be inferred from the observed product fraction. Likewise, the free-particle proposal is quantum Monte Carlo-inspired, but the reported sampling targets are classical distributions and do not include nuclear tunnelling or establish physical-time reaction rates.

Each system was searched once, so these comparisons describe individual search outcomes rather than channel-discovery probabilities. The molecular results measure localization on GFN2-xTB or ff99SB rather than chemical accuracy against experiment. Requiring energies alone extends the choice of energy oracle, although its computational benefit depends on the relative cost of energies and forces.

The conformational errors suggest collective energy-only proposals as a direction for extending continuation to flexible molecules. More generally, the results distinguish finding a connected route from resolving its barrier and characterizing its finite-temperature transition region. This distinction provides a basis for interpreting candidate pathways without conflating geometric agreement with statistical accuracy.

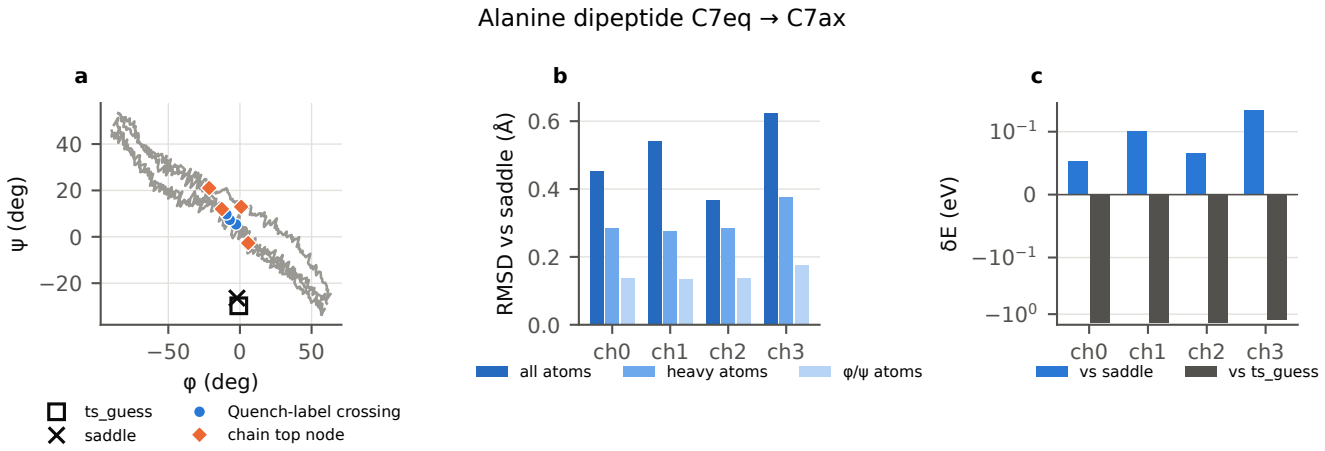


FIG. 3. **Residual conformational error in alanine dipeptide.** **a**, Four chains projected onto the backbone dihedrals (ϕ, ψ). Blue circles mark quench-label crossings, orange diamonds the highest nodes, the black cross the refined ff99SB saddle and the open square the unrelaxed transition-state guess. **b**, Crossing RMSDs to the refined saddle, calculated for all atoms, heavy atoms and the atoms defining ϕ/ψ. **c**, Signed deviations $\delta E = V(\mathbf{R}_0^{\ddagger}) - E_{\text{reference}}$ relative to the refined saddle (blue) and the unrelaxed guess (grey), on a symmetric-logarithmic scale. All four chains came from one search. The positive saddle-referenced differences and displaced ψ values identify error obscured by comparison with the unrelaxed guess.

4. METHODS

tReax_expedition takes two endpoint structures and an energy oracle. Supplementary Algorithm 1 connects exploration, state assignment, chain construction and the optional committor and ensemble calculations. Sections 4.1–4.5 define the statistical targets and reported quantities. The execution details are concentrated in Supplementary Algorithms 2 and 3 and Supplementary Tables S.2–S.4. The validity domains of the resulting quantities are distinguished in Supplementary Table S.1.

4.1. Energy-only sampling in imaginary time

A configuration of N atoms with masses m_i is written $\mathbf{R} = (\mathbf{r}_1, \ldots, \mathbf{r}_N) \in \mathbb{R}^{3N}$, and $V(\mathbf{R})$ denotes its potential energy. The oracle returns an energy or a failed evaluation. Reported search cost counts oracle calls requested by the algorithm, including failed evaluations. Additional evaluations of pre-screened trials in the two-dimensional models are used only for diagnostics and excluded from this count (Supplementary Section 9).

Structures are compared through the mass-weighted squared distance

$$d^2(\mathbf{R}, \mathbf{R}') = \min_{\mathsf{Q} \in SO(3),\, \mathbf{t}} \sum_{i=1}^{N} m_i \left\| \mathbf{r}_i - \mathsf{Q}\mathbf{r}'_i - \mathbf{t} \right\|^2, \qquad (1)$$

where Q is a proper rotation and $\mathbf{t}$ a translation. The minimum is found by centring both structures and applying the mass-weighted Kabsch rotation with the sign correction that excludes reflections[37, 38]. d^2 is symmetric and invariant under a separate rigid motion of either structure (Supplementary Section 2.2). No frame is removed for periodic cells, systems with fixed atoms or the two-dimensional models.

The move is adapted from the Monte Carlo transition-state search (MCTSS) of Tikhonov and Santra[6], which obtains it from real-space diffusion Monte Carlo[30] in imaginary time τ (Supplementary Section 1.2). A walker at parameter τ targets

$$\pi_\tau(\mathbf{R}) = \frac{\exp\left[-\beta(\tau)\, V(\mathbf{R})\right]}{\displaystyle\int \exp\left[-\beta(\tau)\, V(\mathbf{R}')\right] \mathrm{d}\mathbf{R}'}, \qquad (2)$$

where $\beta(\tau) = \tau/\hbar$ is an inverse effective temperature. The integral is over the chosen configuration domain and assumes a finite partition function. For free molecules, translations are removed; dissociating fragments require confinement for an equilibrium interpretation (Supplementary Section 2.4). The integral is never evaluated, because the acceptance test uses only ratios of π_τ. Here τ sets the effective thermal energy $k_{\mathrm{B}}T_{\mathrm{eff}} = \hbar/\tau$, 0.66 eV at $\tau = 1$ fs. It is a sampling parameter rather than a physical time, and equation (2) is a classical Boltzmann factor. The physical temperature corresponds to $\tau^* = \hbar/(k_{\mathrm{B}}T_{\mathrm{phys}})$, or 25.6 fs at 298.15 K.

A trial displaces one uniformly chosen atom i by a Gaussian vector of covariance $\sigma_i^2 \mathsf{I}_3$, with $\sigma_i = \sqrt{\hbar\tau/(2m_i)}$ as in the free-particle kernel of MCTSS[6]. For a single walker the trial is then re-centred and rotated onto the current structure by the proper Kabsch rotation. Every factor of the acceptance probability is invariant under proper rigid motions, so this frame removal leaves the sampled distribution of every invariant quantity exactly unchanged (Supplementary Section 2.4). MCTSS also

places the kinetic factor in the acceptance probability[6]. We use it only as the proposal, because the kinetic factor in the acceptance breaks detailed balance with respect to equation (2) (Supplementary Section 3.3).

Acceptance follows a delayed-acceptance Metropolis–Hastings rule[39–41]. For a move $\mathbf{R} \to \mathbf{R}'$ the acceptance probability is

$$\alpha = \min\bigl(1, e^{\Delta c}\bigr)\, \min\bigl(1, e^{-\beta\Delta(\tilde{V}+b)}\bigr) \times \min\bigl(1, e^{-\beta\Delta(V-\tilde{V})}\bigr), \tag{3}$$

where $\Delta f = f(\mathbf{R}') - f(\mathbf{R})$ for any function f. Here c is the logarithm of a geometric coupling factor that belongs to the target, b an optional bias and $\tilde{V}$ an optional surrogate energy. The three factors are tested in turn, each with its own uniform random number, and the first rejection ends the move. The oracle is therefore called only after the cheap tests pass. The surrogate, GFN2-xTB[34] screening ORCA energies, is used only where stated. Without a surrogate the last two factors merge into $\min(1, e^{-\beta\Delta(V+b)})$. The product of the three ratios equals the Metropolis–Hastings ratio of the target

$$\pi(\mathbf{R}) = \frac{\exp\bigl[c(\mathbf{R}) - \beta V(\mathbf{R}) - \beta b(\mathbf{R})\bigr]}{\displaystyle\int \exp\bigl[c(\mathbf{R}') - \beta V(\mathbf{R}') - \beta b(\mathbf{R}')\bigr]\, \mathrm{d}\mathbf{R}'}, \tag{4}$$

where the denominator normalizes π over all configurations. The move therefore satisfies detailed balance for fixed c, b and $\tilde{V}$, whatever the surrogate (Supplementary Section 3.1). A failed oracle call rejects the trial, as if $V = +\infty$.

4.2. Bidirectional exploration with coupled walkers

Exploration starts two walkers $\mathbf{R}_1$ and $\mathbf{R}_2$ at the endpoints, relaxed by the quench of Section 4.3. It updates them alternately with the move of Section 4.1 on the pair target

$$\Pi(\mathbf{R}_1, \mathbf{R}_2) = \frac{1}{Z_\Pi} \exp\bigl\{-\beta_\mathrm{b}\bigl[V(\mathbf{R}_1) + V(\mathbf{R}_2)\bigr]\bigr\} \times \exp\Bigl[-\frac{d_\mathrm{c}^2(\mathbf{R}_1, \mathbf{R}_2)}{\hbar\tau_{12}}\Bigr], \tag{5}$$

where Z_Π is the integral of the numerator over $\mathbf{R}_1$ and $\mathbf{R}_2$, $\beta_\mathrm{b} = \tau_\mathrm{b}/\hbar$ for the bead parameter τ_b, and $\tau_{12} = n_{12}\tau_\mathrm{b}$ is the coupling parameter. The coupling distance $d_\mathrm{c}^2(\mathbf{R}_1, \mathbf{R}_2) = \sum_i m_i \|\bar{\mathbf{r}}_{1,i} - \bar{\mathbf{r}}_{2,i}\|^2$ is taken between the centred walkers $\bar{\mathbf{R}}_1$ and $\bar{\mathbf{R}}_2$ in their common frame, without rotating one onto the other. Π is therefore invariant under a common rigid motion of both walkers and under a separate translation of either, but not under a separate rotation. Trials are centred and never rotated. For free molecules the relative orientation is instead sampled once per sweep by a rigid rotation of $\mathbf{R}_2$ about a uniformly drawn axis, by a Gaussian angle of standard deviation $\min\{0.5, (\hbar\tau_{12}/\bar{I})^{1/2}\}$ rad. Here $\bar{I}$ is the mean principal moment of inertia of $\mathbf{R}_2$. The rotation leaves V unchanged, so it is accepted on the coupling alone and costs no oracle call. A coupling on the rotation-minimized d^2 would make Π invariant under a separate rotation of either walker and would spoil the midpoint identity below (Supplementary Section 2.5).

MCTSS uses the same free-particle factor to preselect trials that move towards the other walker and redraws a trial until it passes[6]. In tReax_expedition each trial is tested once, as the first factor of equation (3) with $\Delta c = -[d_\mathrm{c}^2(\mathbf{R}_1', \mathbf{R}_2) - d_\mathrm{c}^2(\mathbf{R}_1, \mathbf{R}_2)]/(\hbar\tau_{12})$, and a failed test leaves the walker in place. The coupling thus becomes part of the target. Redrawing until a trial passes would instead tilt the stationary distribution by the pass probability (Supplementary Section 3.2).

Exploration stops at a join and therefore does not provide an equilibrium sample of Π. For a stationary pair, define the midpoint $\mathbf{M} = (\bar{\mathbf{R}}_1 + \bar{\mathbf{R}}_2)/2$ and separation $\mathbf{D} = \bar{\mathbf{R}}_1 - \bar{\mathbf{R}}_2$. The weight

$$w = \exp\bigl\{\beta_\mathrm{b}\bigl[V(\mathbf{R}_1) + V(\mathbf{R}_2) - 2V(\mathbf{M})\bigr]\bigr\} \tag{6}$$

gives a midpoint marginal proportional to $\exp[-2\beta_\mathrm{b}V(\mathbf{M})]$, provided the normalization exists and the relative-coordinate integration domain is independent of the midpoint. The latter holds on the full centred linear configuration space; hard domain restrictions can introduce an additional support factor. The identity follows from a linear change of variables, not a harmonic approximation (Supplementary Section 2.5). Thus $\tau_\mathrm{mid} = 2\tau_\mathrm{b}$. It does not make stopped exploration trajectories equilibrium samples; Supplementary Table S.1 distinguishes these uses.

A run ends at the MCTSS meeting criterion[6], $d^2(\mathbf{R}_1, \mathbf{R}_2)/(\hbar\tau_\mathrm{b}) \le \frac{3}{2}N s_\mathrm{meet}$ with $s_\mathrm{meet} = 1$ for molecules. This test uses the rotation-minimized d^2 of equation (1), because the walkers meet when their shapes do. A run also ends when the two walkers overlap along the mass-weighted axis between the endpoints, or after 5000 steps. The first two outcomes count as a join. The path consists of the walker-1 configurations followed by the walker-2 configurations in reverse order, written in one frame by a common rotation of both walkers.

The bead parameter is chosen from join outcomes on a ladder that first searches for a successful rung and then doubles τ_b until failure. Supplementary Algorithm 1 gives the stopping logic; Supplementary Table S.2 specifies its run budgets and thresholds. The coldest successful rung defines τ_x, the lowest effective temperature at which the walkers join within the search budget. This heuristic depends only on the observed join outcomes. τ_x also fixes subsequent proposal widths and $\ell = \sqrt{3\hbar\tau_\mathrm{x}/2}$, the root-mean-square mass-weighted length of one single-atom move. The resulting joined paths enter the state and channel analysis of Section 4.3.

4.3. States, candidate channels and sampled energy maxima

States are assigned by quenching with the single-walker move of Section 4.1 over decreasing effective temperatures. Proposal widths adapt to acceptance, so a quench is a stochastic optimizer rather than an equilibrium sampler. Supplementary Table S.3 separates the schedules for relaxation, labelling and chain construction. Their budgets and temperature scaling are defined in Supplementary Section 5.4.

The quenched configuration is assigned a state by its fingerprint. For chemical reactions this is the bond graph, with atoms i and j bonded when $r_{ij} < r_i^{\mathrm{cov}} + r_j^{\mathrm{cov}} + 0.4$ Å. For alanine dipeptide it is the backbone dihedral pair (ϕ, ψ) matched within 60°, and for the two-dimensional models the position of the minimum. All fingerprints are invariant under proper rigid motions (Supplementary Section 2.2). The endpoints define states A and B, a new fingerprint registers an intermediate I_k, and every state keeps the lowest minimum found for it (Supplementary Section 5.5).

Up to eight joined exploration paths (twelve for the two-dimensional models), taken from rungs with $\tau_{\mathrm{b}} \leq \tau_{\mathrm{x}}$, are labelled by bisection. The middle frame of every interval whose end labels differ is quenched until each label change lies between adjacent frames. Each such pair of frames is a transition between two states. An excursion that leaves and re-enters a state between two quenched frames is missed, which can lose a channel but cannot create a false one.

Each connected pair of states (S_a, S_b) receives a progress coordinate $\xi(\mathbf{R})$ built from the two fingerprints. For bond graphs, ξ is the summed length of the bonds that break minus that of the bonds that form. For dihedral fingerprints, ξ projects the wrapped dihedral change onto the change between the two minima. Otherwise ξ is the latent output of a classifier trained on the path labels (Supplementary Section 5.6). Chains are grown from up to four transition frames (eight in two dimensions), chosen by farthest-point selection in d.

The interval of ξ between the two minima is divided into slices of width $\Delta\xi$. Node k is obtained by quenching under a harmonic restraint, as in umbrella sampling[42],

$$b_k(\mathbf{R}) = \tfrac{1}{2}\kappa\big[\xi(\mathbf{R}) - \xi_k\big]^2, \tag{7}$$

where ξ_k is the slice centre and $\kappa = \kappa_0/\Delta\xi^2$. Each node starts from its predecessor. A gradual restraint ramp and, when needed, restraint stiffening keep it near its target slice. Large gaps receive intermediate nodes. Unrestrained end quenches identify the states reached by the chain. Supplementary Algorithm 2 specifies this sequence; Supplementary Table S.2 fixes its resolution and tolerances.

Consecutive nodes are joined by two connected discrete paths. On the straight link of mass-weighted length L, V is evaluated at $\max(4, \lceil 4L/\ell \rceil)$ interior points. The quench trajectory, the sequence of accepted configurations that grew a node from its parent, supplies the second path at no extra cost. Keeping the lower maximum of each segment gives the sampled chain maximum

$$E_{\mathrm{up}} = \max\Big\{V(\mathbf{R}_{\mathrm{seed}}),\ \max_k \min\big(E_k^{\mathrm{link}}, E_k^{\mathrm{traj}}\big),\ E_a^{\mathrm{desc}},\ E_b^{\mathrm{desc}}\Big\}, \tag{8}$$

where $\mathbf{R}_{\mathrm{seed}}$ is the first node, E_k^{link} and E_k^{traj} are the largest energies on the link and quench trajectory of segment k, and E_a^{desc} and E_b^{desc} are the largest energies encountered during the two end quenches. A chain is certified when these quenches reach S_a and S_b. E_{up} is then the largest energy on a connected sequence of evaluated configurations between the two basins. If V is L_V-Lipschitz on the interpolating segments, the continuous minimax energy is at most $E_{\mathrm{up}} + \frac{1}{2}L_V h$. Here h includes accepted displacements before frame removal and link-evaluation spacings (Supplementary Section 4.2). Neither L_V nor this correction is computed. Certification therefore concerns endpoint connectivity and evaluated energies, rather than an unconditional continuous-path bound (Supplementary Table S.1). Barriers are reported as $\Delta E_{\mathrm{up}} = E_{\mathrm{up}} - E(S_a)$ relative to the lowest minimum found for S_a.

Certified chains whose highest nodes lie within 3ℓ of each other's nodes form one channel, represented by its lowest chain. On the graph of states weighted by these bounds, the route from A to B minimizes the largest edge value[43]. The route value bounds the minimax energy between A and B only if the minima of each intermediate state that the adjacent chains reach are connected below it (Supplementary Section 4.3). If committor learning is requested and no chain of a pair is certified, the committor is learned from the transition frames first and the chains are regrown along its latent z_θ.

4.4. Quench-label crossing and local ensemble

The majority-quench label of Section 4.3 is bisected along the lowest chain of each channel, first over its nodes and then four times along the straight link that holds the change. Each label is the majority of three quenches with the label schedule of Supplementary Table S.3, to reduce sensitivity to an individual stochastic quench near the ridge. The midpoint of the final interval is the crossing $\mathbf{R}_0^{\ddagger}$, with energy $E_0^{\ddagger}$. It is the midpoint of an interval with opposing majority-quench labels, a finite-budget stochastic construction that need not coincide with the isocommittor surface $q = 1/2$ at T_{phys}. The energy curvature along the chain tangent at the highest node gives a harmonic frequency ω and the crossover temperature $T_{\mathrm{c}} = \hbar|\omega|/(2\pi k_{\mathrm{B}})$ as diagnostics (Supplementary Section 5.8).

On request, a local straddling-pair ensemble (TSE) at T_{phys} is sampled by a pair that straddles the surface $\xi = \xi(\mathbf{R}_0^{\ddagger})$. The pair follows equation (5) of Section 4.2, including the rotation of its relative orientation, with $\tau_{\mathrm{b}} = \tau^*/2$ and $\tau_{12} = 0.05\,\tau_{\mathrm{b}}$. Each walker is confined to its own side of the surface. With the weights of equation (6), the midpoint targets $\exp(-\beta^* V)$, where $\beta^* = \tau^*/\hbar$, within a window around the surface whose width is set by the coupling. For stiff coupling and locally linear ξ, this becomes the Boltzmann distribution on the dividing surface with a surface element weighted by the mass-metric norm of $\nabla\xi$ (Supplementary Section 5.9). Every ensemble reports the Kish effective sample size of its weights[44]. This diagnoses weight concentration and does not account for autocorrelation.

4.5. Committor learning at the physical temperature

Committor learning uses the states, progress coordinates and chains of Section 4.3. For a pair (S_a, S_b), the core C_a contains configurations with the fingerprint of S_a and normalized progress $u = (\xi - \xi_a)/(\xi_b - \xi_a) \le 0.2$, and C_b those of S_b with $u \ge 0.8$. The committor is the first-hitting probability

$$q(\mathbf{R}) = \Pr_{\mathbf{R}}\big[t_{C_b} < t_{C_a}\big], \tag{9}$$

where t_C is the first step at which an unbiased chain of the move of Section 4.1 at τ^*, started at $\mathbf{R}$, enters C[3, 4, 17]. It is therefore the committor of the Monte Carlo chain, not of Newtonian dynamics. For small steps, the move approximates overdamped Langevin dynamics with diffusion constants $D_i = \hbar\tau^*/(4Nm_i)$ per step[45] (Supplementary Section 3.4). The committor of this diffusion minimizes

$$K[q] = \big\langle |\nabla_u q|^2 \big\rangle_{\beta^*} \tag{10}$$

under the conditions $q = 0$ on C_a and $q = 1$ on C_b[19, 26], where the average is over $\exp(-\beta^* V)$ and $|\nabla_u q|^2 = \sum_{j=1}^{3N} m_j^{-1} (\partial q/\partial x_j)^2$ is the squared gradient in the inverse mass metric. Here m_j is the mass of Cartesian coordinate x_j, and the inverse masses follow from the diffusion constants (Supplementary Section 3.5). We treat the small-step correspondence as a modelling assumption. Shooting provides a separate diagnostic of equation (9), with initial-core and finite-budget limitations described in Supplementary Section 8.4.

The model is $q_\theta(\mathbf{R}) = \sigma\big(s\, z_\theta(\mathbf{R})\big)$, where σ is the logistic function, $s = 3$ and z_θ is the latent output of a fully connected network, following fully connected neural committor models[22, 23, 26]. The network reads all interatomic distances for molecules and the coordinates in two dimensions (Supplementary Section 8.1), so the molecular q_θ is invariant under proper rigid motions and its Cartesian gradient is rotation-equivariant (Supplementary Section 2.3). The intended full-data training objective is

$$\begin{aligned}\mathcal{L}(\theta) = \sum_n w_n |\nabla_u q_\theta(\mathbf{R}_n)|^2 + \gamma\, \ell_{\mathrm{shot}}(\theta) \\ + \alpha\big[\langle q_\theta^2\rangle_{C_a} + \langle (1-q_\theta)^2\rangle_{C_b}\big],\end{aligned} \tag{11}$$

where w_n are the normalized, capped reweighting estimates used in training, $\alpha = 1$ weights the boundary penalty and γ weights the shooting term ℓ_{shot}. The gradient term and the boundary penalty follow Kang et al.[26]. The term ℓ_{shot} is the binomial negative log-likelihood of shooting outcomes per trajectory. It goes back to the likelihood maximization of Peters and Trout[46] and is used in the same form in artificial intelligence for molecular mechanism discovery (AIMMD)[20] and in the path sampling built on it[21]. Its weight γ equals Γ times the ratio of the gradient term to ℓ_{shot} at the start of each fit, with $\Gamma = 1$ for molecules and $\Gamma = 0$ in two dimensions. With the boundary conditions imposed as a penalty, the minimizer depends on α and $\mathcal{L}$ is not a strict variational bound[47]. Normalizing weights within each minibatch introduces a further estimator approximation (Supplementary Section 8.2). Shooting outcomes provide diagnostics separate from the training loss.

Training data are generated iteratively, following Kang et al.[26] and Trizio et al.[27]. Unbiased walks from both minima at τ^* provide the core configurations and a first data block. In each of six iterations, walkers started from the minima and from the highest node of each channel (Section 4.3) are driven by $V_{\mathrm{bias}} = V_{\mathrm{K}} + V_z$, which enters the move of Section 4.1 as the bias b. The Kolmogorov bias[26, 27] is

$$\begin{aligned}V_{\mathrm{K}} = -\frac{\lambda}{\beta^*}\Big[\ln\big(|\nabla_u z_\theta|^2 + \epsilon\big) - \ln\epsilon \\ + 2\ln(q_\theta + \delta) + 2\ln(1 - q_\theta + \delta)\Big],\end{aligned} \tag{12}$$

where λ sets the bias strength and ϵ, δ regularize the logarithms (Supplementary Table S.4). For vanishing regularizers V_{K} equals $-(\lambda/\beta^*)\ln|\nabla_u q_\theta|^2$ up to a constant, so it concentrates sampling on the transition region (Supplementary Section 6.2). V_z is a well-tempered metadynamics bias[48, 49] along z_θ that drives crossings between the basins. Biasing along a learned coordinate is now common in enhanced sampling[50], and this role is played by on-the-fly probability enhanced sampling[51] in Trizio et al.[27]. The walkers first build V_z. The bias is then frozen and a production block is recorded under it. The code assigns approximate weights using the multistate Bennett acceptance ratio (MBAR)[36], followed by percentile capping (Supplementary Section 6.4). Capping changes the estimator; its bias is not quantified. The move needs only bias values. The search and learning stages use no derivatives of the potential-energy oracle; network derivatives

enter the bias and loss. Independent reference refinement and Hessian calculations use potential-energy derivatives for evaluation only.

After each production block, shooting points are drawn without replacement from the chain nodes and the new block. Candidate n is drawn with probability

$$P_n = \frac{\left[1 + (z_n/0.5)^2\right]^{-1}/\rho(z_n)}{\sum_m \left[1 + (z_m/0.5)^2\right]^{-1}/\rho(z_m)}, \qquad (13)$$

where $z_n = z_\theta(\mathbf{R}_n)$ and ρ is the histogram density of the candidate latents. This is the Lorentzian selection around $q = 1/2$ of AIMMD[20], which concentrates shooting on the transition region as in shooting from the top[52], here flattened against the sampling density. Two unbiased chains at τ^* run from each point until they enter a core, following the shooting logic of transition path sampling[4, 21, 53]. Their outcomes score the current model before refitting (Supplementary Section 8.4). Finally, 20 chains from each crossing $\mathbf{R}_0^\ddagger$ give the complete-case fraction $p_B = n_B/(n_A + n_B)$ and its binomial standard error for comparison with $q_\theta(\mathbf{R}_0^\ddagger)$. Supplementary Algorithm 3 gives the iterative loop, Supplementary Table S.4 lists its defaults, and Supplementary Table S.6 records system-dependent descriptor sizes and shot budgets.

4.6. Computational details

The workstation had an Intel Xeon 6520P CPU, 1.2 TB of memory and two NVIDIA RTX PRO 6000 Blackwell GPUs, each with 96 GB of VRAM. tReax_expedition is written in Python 3.12 with NumPy[54] and SciPy[55]. Structures were handled with ASE[56], GFN2-xTB energies were evaluated with tblite, and HF-3c energies with ORCA[57]. Networks were built in PyTorch[58] in double precision and trained with Adam[59]. Default settings, random-seed conventions and the software environment are given in Supplementary Sections 7 and 9.9 and Supplementary Tables S.2–S.4.

DATA AVAILABILITY

Benchmark models and starting structures were obtained from three sources. The two-dimensional proton-transfer model and the $Cl^- + CH_3Br \rightarrow CH_3Cl + Br^-$ substitution inputs were adapted from the Supplementary Material of Tikhonov and Santra[6]. The Müller–Brown and two-path potentials and alanine dipeptide inputs were based on the models and data of Trizio, Kang and Parrinello[27], whose supporting data are available in their published repository[60]. Input geometries for the eight remaining gas-phase reactions were obtained from the SACTS benchmark of Suresh and Vogiatzis[33], available at https://github.com/usuresh2509/SACTS.

CODE AVAILABILITY

The tReax_expedition implementation is at https://github.com/FrancisPhillipsBuger/tReax_expedition.

[1] Laidler, K. J. & King, M. C. Development of transition-state theory. *J. Phys. Chem.* **87**, 2657–2664 (1983).

[2] Truhlar, D. G., Garrett, B. C. & Klippenstein, S. J. Current status of transition-state theory. *J. Phys. Chem.* **100**, 12771–12800 (1996).

[3] E, W. & Vanden-Eijnden, E. Transition-path theory and path-finding algorithms for the study of rare events. *Annu. Rev. Phys. Chem.* **61**, 391–420 (2010).

[4] Bolhuis, P. G., Chandler, D., Dellago, C. & Geissler, P. L. Transition path sampling: throwing ropes over rough mountain passes, in the dark. *Annu. Rev. Phys. Chem.* **53**, 291–318 (2002).

[5] Hénin, J., Lelièvre, T., Shirts, M. R., Valsson, O. & Delemotte, L. Enhanced sampling methods for molecular dynamics simulations [article v1.0]. *Living J. Comput. Mol. Sci.* **4**, 1583 (2022).

[6] Tikhonov, D. S. & Santra, R. Two-state reaction path search using a quantum Monte Carlo-inspired approach. *J. Chem. Phys.* **164**, 014114 (2026).

[7] Henkelman, G. & Jónsson, H. A dimer method for finding saddle points on high dimensional potential surfaces using only first derivatives. *J. Chem. Phys.* **111**, 7010–7022 (1999).

[8] Zimmerman, P. M. Single-ended transition state finding with the growing string method. *J. Comput. Chem.* **36**, 601–611 (2015).

[9] Henkelman, G., Uberuaga, B. P. & Jónsson, H. A climbing image nudged elastic band method for finding saddle points and minimum energy paths. *J. Chem. Phys.* **113**, 9901–9904 (2000).

[10] Ásgeirsson, V. *et al.* Nudged elastic band method for molecular reactions using energy-weighted springs combined with eigenvector following. *J. Chem. Theory Comput.* **17**, 4929–4945 (2021).

[11] E, W., Ren, W. & Vanden-Eijnden, E. String method for the study of rare events. *Phys. Rev. B* **66**, 052301 (2002).

[12] Peters, B., Heyden, A., Bell, A. T. & Chakraborty, A. A growing string method for determining transition states: comparison to the nudged elastic band and string methods. *J. Chem. Phys.* **120**, 7877–7886 (2004).

[13] Duan, C., Du, Y., Jia, H. & Kulik, H. J. Accurate transition state generation with an object-aware equivariant elementary reaction diffusion model. *Nat. Comput. Sci.*

3, 1045–1055 (2023).

[14] Duan, C. *et al.* Optimal transport for generating transition states in chemical reactions. *Nat. Mach. Intell.* **7**, 615–626 (2025).

[15] Kim, S., Woo, J. & Kim, W. Y. Diffusion-based generative AI for exploring transition states from 2D molecular graphs. *Nat. Commun.* **15**, 341 (2024).

[16] Lim, H. *et al.* A priori sampling of transition states with guided diffusion. *Preprint at https://arxiv.org/abs/2603.25980* (2026).

[17] Du, R., Pande, V. S., Grosberg, A. Y., Tanaka, T. & Shakhnovich, E. S. On the transition coordinate for protein folding. *J. Chem. Phys.* **108**, 334–350 (1998).

[18] Berezhkovskii, A. & Szabo, A. One-dimensional reaction coordinates for diffusive activated rate processes in many dimensions. *J. Chem. Phys.* **122**, 014503 (2005).

[19] E, W. & Vanden-Eijnden, E. Towards a theory of transition paths. *J. Stat. Phys.* **123**, 503–523 (2006).

[20] Jung, H. *et al.* Machine-guided path sampling to discover mechanisms of molecular self-organization. *Nat. Comput. Sci.* **3**, 334–345 (2023).

[21] Lazzeri, G., Jung, H., Bolhuis, P. G. & Covino, R. Molecular free energies, rates, and mechanisms from data-efficient path sampling simulations. *J. Chem. Theory Comput.* **19**, 9060–9076 (2023).

[22] Khoo, Y., Lu, J. & Ying, L. Solving for high-dimensional committor functions using artificial neural networks. *Res. Math. Sci.* **6**, 1 (2019).

[23] Li, Q., Lin, B. & Ren, W. Computing committor functions for the study of rare events using deep learning. *J. Chem. Phys.* **151**, 054112 (2019).

[24] Megías, A. *et al.* Iterative variational learning of committor-consistent transition pathways using artificial neural networks. *Nat. Comput. Sci.* **5**, 592–602 (2025).

[25] Contreras Arredondo, S. *et al.* Learning the committor without collective variables. *Nat. Comput. Sci.* **6**, 350–357 (2026).

[26] Kang, P., Trizio, E. & Parrinello, M. Computing the committor with the committor to study the transition state ensemble. *Nat. Comput. Sci.* **4**, 451–460 (2024).

[27] Trizio, E., Kang, P. & Parrinello, M. Everything everywhere all at once: a probability-based enhanced sampling approach to rare events. *Nat. Comput. Sci.* **5**, 582–591 (2025).

[28] Bonati, L., Piccini, G. & Parrinello, M. Deep learning the slow modes for rare events sampling. *Proc. Natl Acad. Sci. USA* **118**, e2113533118 (2021).

[29] Zhang, J. *et al.* Descriptor-free collective variables from geometric graph neural networks. *J. Chem. Theory Comput.* **20**, 10787–10797 (2024).

[30] Austin, B. M., Zubarev, D. Y. & Lester, W. A., Jr. Quantum Monte Carlo and related approaches. *Chem. Rev.* **112**, 263–288 (2012).

[31] Müller, K. & Brown, L. D. Location of saddle points and minimum energy paths by a constrained simplex optimization procedure. *Theor. Chim. Acta* **53**, 75–93 (1979).

[32] Baker, J. & Chan, F. The location of transition states: a comparison of Cartesian, Z-matrix, and natural internal coordinates. *J. Comput. Chem.* **17**, 888–904 (1996).

[33] Suresh, U. & Vogiatzis, K. D. Autonomous transition state search with soft actor-critic reinforcement learning. *Preprint at https://arxiv.org/abs/2607.10854* (2026).

[34] Bannwarth, C., Ehlert, S. & Grimme, S. GFN2-xTB—an accurate and broadly parametrized self-consistent tight-binding quantum chemical method with multipole electrostatics and density-dependent dispersion contributions. *J. Chem. Theory Comput.* **15**, 1652–1671 (2019).

[35] Hornak, V. *et al.* Comparison of multiple Amber force fields and development of improved protein backbone parameters. *Proteins* **65**, 712–725 (2006).

[36] Shirts, M. R. & Chodera, J. D. Statistically optimal analysis of samples from multiple equilibrium states. *J. Chem. Phys.* **129**, 124105 (2008).

[37] Kabsch, W. A solution for the best rotation to relate two sets of vectors. *Acta Crystallogr. A* **32**, 922–923 (1976).

[38] Kabsch, W. A discussion of the solution for the best rotation to relate two sets of vectors. *Acta Crystallogr. A* **34**, 827–828 (1978).

[39] Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H. & Teller, E. Equation of state calculations by fast computing machines. *J. Chem. Phys.* **21**, 1087–1092 (1953).

[40] Hastings, W. K. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **57**, 97–109 (1970).

[41] Christen, J. A. & Fox, C. Markov chain Monte Carlo using an approximation. *J. Comput. Graph. Stat.* **14**, 795–810 (2005).

[42] Torrie, G. M. & Valleau, J. P. Nonphysical sampling distributions in Monte Carlo free-energy estimation: umbrella sampling. *J. Comput. Phys.* **23**, 187–199 (1977).

[43] Pollack, M. The maximum capacity through a network. *Oper. Res.* **8**, 733–736 (1960).

[44] Kish, L. *Survey Sampling* (Wiley, New York, 1965).

[45] Kikuchi, K., Yoshida, M., Maekawa, T. & Watanabe, H. Metropolis Monte Carlo method as a numerical technique to solve the Fokker–Planck equation. *Chem. Phys. Lett.* **185**, 335–338 (1991).

[46] Peters, B. & Trout, B. L. Obtaining reaction coordinates by likelihood maximization. *J. Chem. Phys.* **125**, 054108 (2006).

[47] Petersen, M., Lichtinger, S. & Covino, R. Committors and reaction rates from trial functions that violate the boundary conditions. *Preprint at https://arxiv.org/abs/2608.02536* (2026).

[48] Laio, A. & Parrinello, M. Escaping free-energy minima. *Proc. Natl Acad. Sci. USA* **99**, 12562–12566 (2002).

[49] Barducci, A., Bussi, G. & Parrinello, M. Well-tempered metadynamics: A smoothly converging and tunable free-energy method. *Phys. Rev. Lett.* **100**, 020603 (2008).

[50] Mehdi, S., Smith, Z., Herron, L., Zou, Z. & Tiwary, P. Enhanced sampling with machine learning. *Annu. Rev. Phys. Chem.* **75**, 347–370 (2024).

[51] Invernizzi, M. & Parrinello, M. Rethinking metadynamics: from bias potentials to probability distributions. *J. Phys. Chem. Lett.* **11**, 2731–2736 (2020).

[52] Jung, H., Okazaki, K.-i. & Hummer, G. Transition path sampling of rare events by shooting from the top. *J.*

Chem. Phys. **147**, 152716 (2017).

[53] Dellago, C., Bolhuis, P. G. & Chandler, D. Efficient transition path sampling: application to Lennard-Jones cluster rearrangements. *J. Chem. Phys.* **108**, 9236–9245 (1998).

[54] Harris, C. R. *et al.* Array programming with NumPy. *Nature* **585**, 357–362 (2020).

[55] Virtanen, P. *et al.* SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods* **17**, 261–272 (2020).

[56] Larsen, A. H. *et al.* The atomic simulation environment—a Python library for working with atoms. *J. Phys. Condens. Matter* **29**, 273002 (2017).

[57] Neese, F., Wennmohs, F., Becker, U. & Riplinger, C. The ORCA quantum chemistry program package. *J. Chem. Phys.* **152**, 224108 (2020).

[58] Paszke, A. *et al.* PyTorch: an imperative style, high-performance deep learning library. *Adv. Neural Inf. Process. Syst.* **32**, 8024–8035 (2019).

[59] Kingma, D. P. & Ba, J. Adam: a method for stochastic optimization. *Preprint at https://arxiv.org/abs/1412.6980* (2014).

[60] Trizio, E., Kang, P. & Parrinello, M. Supporting data for the paper "Everything everywhere all at once, a probability-based enhanced sampling approach to rare events". GitHub https://github.com/EnricoTrizio/committor_2.0 (2025).

[61] Feynman, R. P. & Hibbs, A. R. *Quantum Mechanics and Path Integrals* (McGraw-Hill, New York, 1965).

[62] Kemeny, J. G. & Snell, J. L. *Finite Markov Chains* (Van Nostrand, Princeton, 1960).

[63] Tierney, L. Markov chains for exploring posterior distributions. *Ann. Stat.* **22**, 1701–1728 (1994).

[64] Lelièvre, T. & Stoltz, G. Partial differential equations and stochastic methods in molecular dynamics. *Acta Numer.* **25**, 681–880 (2016).

[65] Bovier, A., Eckhoff, M., Gayrard, V. & Klein, M. Metastability in reversible diffusion processes I: sharp asymptotics for capacities and exit times. *J. Eur. Math. Soc.* **6**, 399–424 (2004).

[66] Gneiting, T. & Raftery, A. E. Strictly proper scoring rules, prediction, and estimation. *J. Am. Stat. Assoc.* **102**, 359–378 (2007).

Supplementary Information for Candidate reaction pathways from energy-only Monte Carlo exploration and continuation

Chien-Chai Chang
Department of Engineering and System Science,
National Tsing Hua University (NTHU), Hsinchu City, 300044, Taiwan

1. BACKGROUND

1.1. Notation, units and effective temperatures

Energies are in eV, lengths in Å, masses in atomic mass units (amu) and the time parameters τ in fs. The inverse effective temperature of one walker is $\beta(\tau) = \tau/\hbar$, so $k_{\mathrm{B}}T_{\mathrm{eff}} = \hbar/\tau$. With $\hbar = 0.6582$ eV fs, $\tau = 1$ fs corresponds to 0.658 eV (63.5 kJ mol^{-1}), and each doubling of τ halves this value. The physical temperature enters only through $\tau^* = \hbar/(k_{\mathrm{B}}T_{\mathrm{phys}})$, which is 25.62 fs at 298.15 K and 25.46 fs at 300 K. For the two-dimensional models in reduced units, one energy unit is $k_{\mathrm{B}} \times 300$ K and $\tau^* = 25.46$ fs. For the proton-transfer model, whose energies are in cm^{-1}, $\tau^* = 10$ fs places the barrier near $8\,k_{\mathrm{B}}T^*$.

The Gaussian proposal of the kernel has standard deviation $\sigma_i = \lambda_{\mathrm{s}}\sqrt{\hbar\tau/(2m_i)}$ per Cartesian component. One single-atom move therefore has mean mass-weighted squared length $3m_i\sigma_i^2 = \frac{3}{2}\hbar\tau\lambda_{\mathrm{s}}^2$, independent of the atom, which defines $\ell = \sqrt{3\hbar\tau_{\mathrm{x}}/2}$ for $\lambda_{\mathrm{s}} = 1$. At $\tau = 1$ fs, $\ell = 0.098$ amu$^{1/2}$ Å.

1.2. Origin of the Monte Carlo move

MCTSS derives its move from real-space diffusion Monte Carlo[6, 30]. An imaginary time step $\Delta t = -\mathrm{i}\tau$ and a second-order Trotter splitting of $\exp(-\tau\hat{H}/\hbar)$ give the unnormalized transition weight

$$\tilde{p}(\mathbf{R} \to \mathbf{R}') = \exp\Big[-\frac{\tau V(\mathbf{R})}{\hbar} - \frac{\tau V(\mathbf{R}')}{\hbar}\Big] \times \exp\Big[-\frac{d^2(\mathbf{R},\mathbf{R}')}{\hbar\tau}\Big], \tag{S.1}$$

where d^2 is the distance of equation (1). The last factor is the squared free-particle propagator[61]. MCTSS accepts $\mathbf{R}'$ with $\min[1, \tilde{p}(\mathbf{R} \to \mathbf{R}')/\tilde{p}(\mathbf{R} \to \mathbf{R})]$ and draws $\mathbf{R}'$ by displacing one atom with a Gaussian of variance $\hbar\tau/(2m_i)$ per component[6]. For a single-atom move of atom i by $\boldsymbol{\delta}$, the density of that proposal is

$$g(\boldsymbol{\delta}) = \Big(\frac{m_i}{\pi\hbar\tau}\Big)^{3/2} \exp\Big[-\frac{m_i\|\boldsymbol{\delta}\|^2}{\hbar\tau}\Big], \tag{S.2}$$

whose exponent is the kinetic factor of equation (S.1). The kinetic factor therefore acts twice. tReax_expedition keeps the Gaussian proposal and removes the kinetic factor from the acceptance, which leaves the Boltzmann factor $\exp[-\beta(\tau)V]$ of equation (2) as the stationary distribution. Supplementary Section 3.3 shows that the kinetic factor in the acceptance breaks detailed balance with respect to this factor. The directional pre-selection of MCTSS, $\exp[-d^2/(\hbar\tau_{12})]$ with $\tau_{12} = N_{\mathrm{inter}}\tau$ in their notation, becomes the coupling factor of the pair target, equation (5). In tReax_expedition this factor uses the common-frame distance d_{c}^2 instead of d^2 (Supplementary Section 2.5). The meeting criterion is taken over unchanged (Supplementary Section 5.2).

1.3. Classes of modelling elements

The method combines exact statistical-mechanical relations with algorithmic devices whose parameters have no physical meaning. Supplementary Table S.1 assigns every element to one class and states the assumption under which it holds. Claims that depend on an assumption in the last column should be read with that assumption in mind.

2. PROOF OF SE(3) INVARIANCE AND EQUIVARIANCE

This section establishes the rigid-motion properties used by the sampling kernel and committor model. For an isolated molecule, the energy and interatomic-distance descriptors are invariant under proper rigid motions, whereas the network gradients entering the Kolmogorov bias and loss are equivariant. These properties imply that removing the frame of a single walker, or retaining a coupled pair in a common frame, preserves the sampled distribution of invariant observables. The derivation below makes the required assumptions explicit; the released unit tests provide numerical checks of the same identities.

2.1. Notation

Let $\mathsf{Q} \in SO(3)$ be a proper rotation matrix and $\mathbf{t} \in \mathbb{R}^3$ a translation vector. The rigid motion $g = (\mathsf{Q}, \mathbf{t})$ acts on a configuration as $g\mathbf{R} = (\mathsf{Q}\mathbf{r}_1 + \mathbf{t}, \ldots, \mathsf{Q}\mathbf{r}_N + \mathbf{t})$, or in $3N$-dimensional notation

$$g\mathbf{R} = \hat{\mathsf{Q}}\mathbf{R} + \mathbf{1}_N \otimes \mathbf{t}, \tag{S.3}$$

where $\hat{\mathsf{Q}} = \mathsf{I}_N \otimes \mathsf{Q}$ is block diagonal and $\mathbf{1}_N$ is the vector of N ones. The mass matrix $\mathsf{W} = \mathrm{diag}(m_1, \ldots, m_N) \otimes \mathsf{I}_3$ multiplies each atomic block by a scalar, so that

$$\hat{\mathsf{Q}}^{\mathsf{T}}\mathsf{W}\hat{\mathsf{Q}} = \mathsf{W}, \qquad \hat{\mathsf{Q}}^{\mathsf{T}}\mathsf{W}^{-1}\hat{\mathsf{Q}} = \mathsf{W}^{-1}. \tag{S.4}$$

TABLE S.1. **Definitions, assumptions and numerical approximations underlying the reported quantities.** The entries distinguish fixed-target identities from modelling assumptions and finite-resolution search procedures. Conditions in the final column delimit each interpretation. In particular, stationary pair identities do not apply automatically to searches stopped at a join, and endpoint certification does not evaluate the continuous-path correction.

Element	Class	Remark
$V(\mathbf{R})$, masses, T_{phys}	Physical definition	Level of theory fixed per system (Supplementary Table S.5).
$\exp(-\beta^* V)$, committor q of eq. (9), transition-state ensemble	Statistical-mechanical construct	q is defined for the Monte Carlo chain at τ^*.
Functional $K[q]$ of eq. (10)	Statistical-mechanical construct	Minimized by q only in the small-step limit (Supplementary Sections 3.4 and 3.5).
$\beta(\tau) = \tau/\hbar$ away from τ^*, τ ladder, pair target Π	Algorithmic device	Effective temperatures of exploration, not physical ones.
Midpoint ensemble at $2\tau_{\mathrm{b}}$	Algorithmic device	Weighted identity for a stationary, normalizable pair target in a common frame (Supplementary Section 2.5).
Quenches, fingerprints, majority-quench labels	Algorithmic device	Basins defined by a stochastic optimizer.
Restrained chains, E_{up}, channels, minimax route	Algorithmic device	Endpoint connectivity checked; continuous bound needs the uncomputed $\frac{1}{2}L_V h$ correction (Supplementary Section 4.2). Routes also require connected intermediate states.
Quench-label crossing $\mathbf{R}_0^{\ddagger}$	Algorithmic device	Finite-budget label-change estimate, not a verified separatrix or $q = 1/2$.
Frame removal of a single walker, common frame of a pair	Exact algorithmic device	Exact for invariant observables (Supplementary Sections 2.4 and 2.5).
Surface measure of the transition-state ensemble	Algorithmic approximation	Stiff coupling and locally linear ξ.
q_θ, ξ_{label}, penalty boundary term, shot likelihood	Machine-learning approximation	Finite-capacity model with penalty boundary conditions; assessed by shooting.
MBAR weights	Statistical estimator	Equilibrium estimator for fixed biases and overlapping sampled states; weight capping adds bias.
Weight cap, adaptive widths, meeting and success thresholds, slice number, restraint schedule, densification, channel radius, votes, core fraction, Lorentzian width	Numerical heuristic	Defaults in Supplementary Tables S.2–S.4.

We write $\|\mathbf{x}\|_{\mathsf{W}}^2 = \mathbf{x}^{\mathsf{T}}\mathsf{W}\mathbf{x}$. The centring map is $C(\mathbf{R}) = \mathbf{R} - \mathbf{1}_N \otimes \bar{\mathbf{r}}(\mathbf{R})$, where $\bar{\mathbf{r}} = \sum_i m_i \mathbf{r}_i / \sum_i m_i$ is the centre of mass. Since $\bar{\mathbf{r}}(g\mathbf{R}) = \mathsf{Q}\bar{\mathbf{r}}(\mathbf{R}) + \mathbf{t}$,

$$C(g\mathbf{R}) = \hat{\mathsf{Q}}\, C(\mathbf{R}). \tag{S.5}$$

For an isolated molecule we require of the oracle and of any surrogate that

$$V(g\mathbf{R}) = V(\mathbf{R}), \qquad \tilde{V}(g\mathbf{R}) = \tilde{V}(\mathbf{R}), \tag{S.6}$$

for all g. GFN2-xTB, HF-3c and the force field of alanine dipeptide satisfy equation (S.6). Frame removal is switched off for periodic cells, for systems with fixed atoms and for the two-dimensional models, where equation (S.6) is not needed.

2.2. Invariance of distances, fingerprints and the committor model

We prove that every scalar quantity used by the method is invariant to rigid motions.

Rotation-minimized distance. The optimal transla-

tion in equation (1) superposes the two centres of mass, so that

$$d^2(\mathbf{R},\mathbf{R}') = \min_{\mathsf{Q}''\in SO(3)} \big\|C(\mathbf{R}) - \hat{\mathsf{Q}}''C(\mathbf{R}')\big\|_{\mathsf{W}}^2. \tag{S.7}$$

For two rigid motions $g=(\mathsf{Q},\mathbf{t})$ and $g'=(\mathsf{Q}',\mathbf{t}')$, equations (S.4) and (S.5) give

$$\begin{aligned}\big\|C(g\mathbf{R}) - \hat{\mathsf{Q}}''C(g'\mathbf{R}')\big\|_{\mathsf{W}} &= \big\|\hat{\mathsf{Q}}C(\mathbf{R}) - \hat{\mathsf{Q}}''\hat{\mathsf{Q}}'C(\mathbf{R}')\big\|_{\mathsf{W}} \\ &= \big\|C(\mathbf{R}) - \hat{\mathsf{Q}}^{\mathsf{T}}\hat{\mathsf{Q}}''\hat{\mathsf{Q}}'C(\mathbf{R}')\big\|_{\mathsf{W}}.\end{aligned} \tag{S.8}$$

As Q'' runs over $SO(3)$, so does $\mathsf{Q}^{\mathsf{T}}\mathsf{Q}''\mathsf{Q}'$. The two minima therefore coincide, and $d^2(g\mathbf{R},g'\mathbf{R}') = d^2(\mathbf{R},\mathbf{R}')$ for separate rigid motions of the two structures. Since $\|C(\mathbf{R})-\hat{\mathsf{Q}}''C(\mathbf{R}')\|_{\mathsf{W}} = \|\hat{\mathsf{Q}}''^{\mathsf{T}}C(\mathbf{R})-C(\mathbf{R}')\|_{\mathsf{W}}$ and $\mathsf{Q}''^{\mathsf{T}}\in SO(3)$, d^2 is also symmetric.

Common-frame distance. The coupling of equation (5) uses $d_{\mathrm{c}}^2(\mathbf{R},\mathbf{R}') = \|C(\mathbf{R})-C(\mathbf{R}')\|_{\mathsf{W}}^2$. For a common rigid motion,

$$d_{\mathrm{c}}^2(g\mathbf{R},g\mathbf{R}') = \big\|\hat{\mathsf{Q}}\big[C(\mathbf{R})-C(\mathbf{R}')\big]\big\|_{\mathsf{W}}^2 = d_{\mathrm{c}}^2(\mathbf{R},\mathbf{R}'). \tag{S.9}$$

A translation of one structure does not change its centred coordinates, so d_{c}^2 is also invariant to separate translations. It is not invariant to a rotation of one structure alone, which is why the relative orientation of a pair is sampled by the explicit move of Supplementary Section 2.5.

Interatomic distances and descriptors. For every pair of atoms,

$$\big\|(\mathsf{Q}\mathbf{r}_i+\mathbf{t})-(\mathsf{Q}\mathbf{r}_j+\mathbf{t})\big\| = \big\|\mathsf{Q}(\mathbf{r}_i-\mathbf{r}_j)\big\| = \|\mathbf{r}_i-\mathbf{r}_j\|. \tag{S.10}$$

The bond graph compares r_{ij} with fixed thresholds, ξ_{bond} is a signed sum of r_{ij}, and the descriptors are $(r_{ij}^2+10^{-12})^{1/2}$. The network outputs z_θ and q_θ depend on $\mathbf{R}$ only through the descriptors. Therefore all of them are invariant.

Dihedral angles. For atoms a,b,c,d let $\mathbf{b}_1=\mathbf{r}_b-\mathbf{r}_a$, $\mathbf{b}_2=\mathbf{r}_c-\mathbf{r}_b$, $\mathbf{b}_3=\mathbf{r}_d-\mathbf{r}_c$, $\mathbf{n}_1=\mathbf{b}_1\times\mathbf{b}_2$ and $\mathbf{n}_2=\mathbf{b}_2\times\mathbf{b}_3$. The dihedral angle is

$$\phi = \operatorname{atan2}\Big(\frac{(\mathbf{n}_1\times\mathbf{n}_2)\cdot\mathbf{b}_2}{\|\mathbf{b}_2\|},\ \mathbf{n}_1\cdot\mathbf{n}_2\Big). \tag{S.11}$$

Translations cancel in the $\mathbf{b}_k$. For $\det\mathsf{Q}=+1$, $\mathsf{Q}(\mathbf{u}\times\mathbf{v}) = (\mathsf{Q}\mathbf{u})\times(\mathsf{Q}\mathbf{v})$ and scalar products are preserved, so both arguments of atan2 are unchanged. Therefore the dihedral fingerprint and ξ_{dih} are invariant.

Progress along the endpoint axis. The progress $p(\mathbf{R})$ of the overlap criterion (Supplementary Section 5.2) is evaluated after aligning $\mathbf{R}$ onto a fixed endpoint. The aligned structure is the element of the orbit $\{g\mathbf{R}\}$ closest to that endpoint, which is the same for $\mathbf{R}$ and for $g\mathbf{R}$. Therefore p is invariant.

Reflections and permutations. Under an improper rotation, distances and bond graphs are unchanged, whereas dihedral angles change sign. The committor model and the bond-graph fingerprint therefore cannot distinguish mirror images, but the dihedral fingerprint can. The chain never applies a reflection. Moves are displacements, and the Kabsch rotation is restricted to $\det\mathsf{Q}=+1$ by its sign correction[38]. Atoms carry fixed labels, and configurations that differ by a permutation of identical atoms have different labelled bond graphs. They are registered as different states. This does not affect the statements in this section, but symmetry-equivalent channels are reported separately.

2.3. Equivariance of gradients and invariance of biases

We now prove that the gradient of an invariant function is equivariant to rotations and that its norm in the inverse mass metric is invariant.

Equivariance. Let $f:\mathbb{R}^{3N}\to\mathbb{R}$ be differentiable with $f(g\mathbf{R})=f(\mathbf{R})$ for all g. Differentiating $f(\hat{\mathsf{Q}}\mathbf{R}+\mathbf{1}_N\otimes\mathbf{t}) = f(\mathbf{R})$ with respect to $\mathbf{R}$ gives $\hat{\mathsf{Q}}^{\mathsf{T}}\nabla f(g\mathbf{R}) = \nabla f(\mathbf{R})$, and multiplication by $\hat{\mathsf{Q}}$ gives

$$\nabla f(g\mathbf{R}) = \hat{\mathsf{Q}}\,\nabla f(\mathbf{R}). \tag{S.12}$$

Invariance of the gradient norm. With equation (S.4), the squared gradient in the inverse mass metric, $|\nabla_u f|^2 = \nabla f^{\mathsf{T}}\mathsf{W}^{-1}\nabla f$, satisfies

$$\begin{aligned}|\nabla_u f|^2(g\mathbf{R}) &= \nabla f(\mathbf{R})^{\mathsf{T}}\hat{\mathsf{Q}}^{\mathsf{T}}\mathsf{W}^{-1}\hat{\mathsf{Q}}\,\nabla f(\mathbf{R}) \\ &= |\nabla_u f|^2(\mathbf{R}).\end{aligned} \tag{S.13}$$

Biases and targets. Taking $f=z_\theta$, which is invariant by Supplementary Section 2.2, shows that $|\nabla_u z_\theta|^2$ is invariant. The Kolmogorov bias of equation (12) depends on $\mathbf{R}$ only through $|\nabla_u z_\theta|^2$ and q_θ, the adaptive bias V_z only through z_θ, and the restraints b_k and the side constraint of the transition-state ensemble only through ξ. Together with equation (S.6), this gives

$$\begin{aligned}\pi(g\mathbf{R}) &= \pi(\mathbf{R}),\\ \Pi(g\mathbf{R}_1,g\mathbf{R}_2) &= \Pi(\mathbf{R}_1,\mathbf{R}_2),\\ \Pi(\mathbf{R}_1+\mathbf{1}_N\otimes\mathbf{t}_1,\mathbf{R}_2+\mathbf{1}_N\otimes\mathbf{t}_2) &= \Pi(\mathbf{R}_1,\mathbf{R}_2),\end{aligned} \tag{S.14}$$

for the targets of equations (4) and (5). Therefore, we can conclude that every target of the method is invariant to rigid motions, and Π is invariant to a common rigid motion of both walkers and to separate translations. The same argument shows that the loss of equation (11) does not change when any training configuration is moved rigidly, so augmenting the data with rotated copies would add no information.

2.4. Exactness of frame removal for a single walker

We prove that aligning each trial onto the current structure leaves the distribution of every invariant quantity unchanged. Let P be the kernel of the move of Methods Section 4.1 without frame removal. It selects atom i with probability $1/N$, proposes $\mathbf{y} = \mathbf{R} + \mathbf{e}_i \otimes \boldsymbol{\delta}$ with $\boldsymbol{\delta} \sim \mathcal{N}(\mathbf{0}, \sigma_i^2 \mathsf{I}_3)$, where $\mathbf{e}_i$ is the i-th unit vector of $\mathbb{R}^N$, and accepts with $\alpha(\mathbf{R}, \mathbf{y})$ of equation (3). Let P_{A} be the implemented kernel, in which the trial is replaced by $A_{\mathbf{R}}(\mathbf{y}) = \hat{\mathsf{Q}}^* C(\mathbf{y})$, where $\mathsf{Q}^* \in SO(3)$ is the Kabsch rotation of $C(\mathbf{y})$ onto $C(\mathbf{R})$. A set B is called invariant when $gB = B$ for all g. It is a union of orbits $[\mathbf{R}] = \{g\mathbf{R}\}$, and an invariant observable is a function of the orbit alone.

Acceptance. The acceptance probability depends on the trial only through c, V, b and $\tilde{V}$, which are invariant (Supplementary Section 2.3). Since $A_{\mathbf{R}}(\mathbf{y})$ lies in the orbit of $\mathbf{y}$,

$$\alpha\big(\mathbf{R}, A_{\mathbf{R}}(\mathbf{y})\big) = \alpha(\mathbf{R}, \mathbf{y}). \tag{S.15}$$

Invariant sets. For invariant B, $A_{\mathbf{R}}(\mathbf{y}) \in B$ if and only if $\mathbf{y} \in B$. A rejected trial leaves $\mathbf{R}$ in place under both kernels, so

$$P_{\mathrm{A}}(\mathbf{R}, B) = \mathbb{E}\big[\alpha \mathbb{1}_B(\mathbf{y})\big] + \big(1 - \mathbb{E}[\alpha]\big)\mathbb{1}_B(\mathbf{R}) = P(\mathbf{R}, B), \tag{S.16}$$

where the expectation is over the atom and the displacement.

Equivariance of the kernel. From $g\mathbf{R}$ the trial is

$$g\mathbf{R} + \mathbf{e}_i \otimes \boldsymbol{\delta} = g\big(\mathbf{R} + \mathbf{e}_i \otimes \mathsf{Q}^{\mathsf{T}}\boldsymbol{\delta}\big). \tag{S.17}$$

The displacement is isotropic, so $\mathsf{Q}^{\mathsf{T}}\boldsymbol{\delta}$ has the same distribution as $\boldsymbol{\delta}$, and the trial from $g\mathbf{R}$ is distributed as g applied to the trial from $\mathbf{R}$. With $\alpha(g\mathbf{R}, g\mathbf{y}) = \alpha(\mathbf{R}, \mathbf{y})$ this gives

$$P(g\mathbf{R}, B) = P(\mathbf{R}, g^{-1}B) = P(\mathbf{R}, B). \tag{S.18}$$

Hence $P(\mathbf{R}, B)$ depends on $\mathbf{R}$ only through its orbit $[\mathbf{R}]$.

Induction over steps. Let the chains $\mathbf{R}_n$ of P and $\mathbf{R}_n^{\mathrm{A}}$ of P_{A} start from the same configuration $\mathbf{R}_0$. We show by induction that their orbits have the same distribution at every step.

- *Base case ($n = 0$):* Both chains start at $\mathbf{R}_0$, so $\Pr\big[\mathbf{R}_0^{\mathrm{A}} \in B\big] = \Pr\big[\mathbf{R}_0 \in B\big]$ for every invariant B.
- *Induction hypothesis:* At step $n = k$, $\Pr\big[\mathbf{R}_k^{\mathrm{A}} \in B\big] = \Pr\big[\mathbf{R}_k \in B\big]$ for every invariant B.
- *Induction step:* By equation (S.16),

$$\Pr\big[\mathbf{R}_{k+1}^{\mathrm{A}} \in B\big] = \mathbb{E}\big[P_{\mathrm{A}}(\mathbf{R}_k^{\mathrm{A}}, B)\big] = \mathbb{E}\big[P(\mathbf{R}_k^{\mathrm{A}}, B)\big]. \tag{S.19}$$

By equation (S.18), $P(\cdot, B)$ is an invariant function, so its expectation depends only on the distribution of the orbit, which is the same for both chains by the induction hypothesis. Hence

$$\Pr\big[\mathbf{R}_{k+1}^{\mathrm{A}} \in B\big] = \mathbb{E}\big[P(\mathbf{R}_k, B)\big] = \Pr\big[\mathbf{R}_{k+1} \in B\big]. \tag{S.20}$$

By induction, the orbits of the two chains have the same distribution at every step. Applying the same argument to joint events at several steps shows that invariant observables also have the same time correlations. The orbit process is a lumped Markov chain[62] with kernel $\bar{P}([\mathbf{R}], B) = P(\mathbf{R}, B)$. P satisfies detailed balance with respect to π (Supplementary Section 3.1), and integrating it over invariant sets B_1 and B_2 gives

$$\int_{B_1} \pi(\mathbf{R}) P(\mathbf{R}, B_2)\, \mathrm{d}\mathbf{R} = \int_{B_2} \pi(\mathbf{R}) P(\mathbf{R}, B_1)\, \mathrm{d}\mathbf{R}, \tag{S.21}$$

which is detailed balance of the lumped chain with respect to the image of π. Frame removal therefore preserves the process of invariant observables.

Remarks. (a) The proof does not require the aligned proposal to be symmetric. Frame removal is therefore exact rather than an approximation. (b) Along translations π is not normalizable, so the statement is applied to centred configurations. The image of π on molecular shapes is then normalizable when $\exp(-\beta V)$ is integrable over shapes. This fails for fragments that separate in free space, as in $\mathrm{H_2CO} \to \mathrm{H_2} + \mathrm{CO}$, a property of the gas-phase ensemble rather than of frame removal. (c) Linear configurations, such as HCN and HCCH near their minima, have lower-dimensional orbits and a Kabsch rotation that is not unique. They form a set of measure zero, and any choice keeps the trial in its orbit. (d) The result applies to every single-walker chain of the method, namely quenches, restrained nodes, biased walks, shots and validation chains. (e) If a surrogate violated equation (S.6), stationarity would still hold (Supplementary Section 3.1), but the acceptance step above would fail.

2.5. Exactness of the common frame for a coupled pair

For the full centred linear configuration space, we prove that the pair sweep leaves Π of equation (5) invariant and that the midpoint weights are exact. One sweep consists of two kinds of move. (a) For $n = 1, 2$, walker n is displaced as in Supplementary Section 2.4, centred without rotation, and tested with equation (3) with $c = -d_{\mathrm{c}}^2/(\hbar\tau_{12})$. (b) The centred walker $\bar{\mathbf{R}}_2$ is rotated to $\hat{\mathsf{Q}}_\omega \bar{\mathbf{R}}_2$, where Q_ω rotates by the angle $\omega \sim \mathcal{N}(0, \sigma_{\mathrm{rot}}^2)$ about an axis drawn uniformly on the unit sphere and σ_{rot} is given by equation (S.48).

Displacements. At fixed $\mathbf{R}_m$, the displacement of walker n with the coupling test is the move of Supplementary Section 3.1 for the conditional of Π, so it satisfies detailed balance with respect to that conditional. Centring is a translation of walker n alone, under which Π is invariant by equation (S.14). The argument of Supple-

mentary Section 2.4, with the translations of walker n in place of $SE(3)$, shows that centring leaves this conclusion unchanged. An update that preserves the conditional of Π preserves Π.

Relative rotation. The inverse of Q_ω is the rotation by $-\omega$ about the same axis, which has the same probability because the distribution of ω is symmetric. σ_{rot} depends on $\bar{\mathbf{R}}_2$ only through the trace of its inertia tensor, which a rotation does not change. The proposal is therefore symmetric with respect to the rotation-invariant measure on the orbit of $\bar{\mathbf{R}}_2$. Because $V(\hat{\mathsf{Q}}_\omega\bar{\mathbf{R}}_2) = V(\bar{\mathbf{R}}_2)$, the Metropolis–Hastings ratio of Π reduces to

$$\frac{\Pi(\mathbf{R}_1, \hat{\mathsf{Q}}_\omega\bar{\mathbf{R}}_2)}{\Pi(\mathbf{R}_1, \bar{\mathbf{R}}_2)} = \exp\Big[-\frac{d_{\mathrm{c}}^2(\mathbf{R}_1, \hat{\mathsf{Q}}_\omega\bar{\mathbf{R}}_2) - d_{\mathrm{c}}^2(\mathbf{R}_1, \bar{\mathbf{R}}_2)}{\hbar\tau_{12}}\Big], \tag{S.22}$$

which is the ratio used in the test. Detailed balance with respect to Π follows. A composition of kernels that each preserve Π preserves Π[63], so the whole sweep does.

Midpoint identity. Let $\mathbf{M} = (\bar{\mathbf{R}}_1 + \bar{\mathbf{R}}_2)/2$ and $\mathbf{D} = \bar{\mathbf{R}}_1 - \bar{\mathbf{R}}_2$. The map $(\bar{\mathbf{R}}_1, \bar{\mathbf{R}}_2) \mapsto (\mathbf{M}, \mathbf{D})$ is linear with a Jacobian of unit modulus, and $d_{\mathrm{c}}^2 = \mathbf{D}^{\mathsf{T}}\mathsf{W}\mathbf{D}$ without any rotation. Multiplying Π by the weight w of equation (6) replaces $\exp\{-\beta_{\mathrm{b}}[V(\mathbf{R}_1) + V(\mathbf{R}_2)]\}$ by $\exp[-2\beta_{\mathrm{b}}V(\mathbf{M})]$, which gives

$$\Pi(\mathbf{R}_1, \mathbf{R}_2)\, w = \frac{1}{Z_\Pi}\exp\big[-2\beta_{\mathrm{b}}V(\mathbf{M})\big]\exp\Big[-\frac{\mathbf{D}^{\mathsf{T}}\mathsf{W}\mathbf{D}}{\hbar\tau_{12}}\Big], \tag{S.23}$$

where Z_Π is the normalization of Π. When the normalization exists, the weighted midpoint marginal is proportional to $\exp[-2\beta_{\mathrm{b}}V(\mathbf{M})]$. This identity presumes stationary sampling of Π and a midpoint-independent integration domain for $\mathbf{D}$. Hard confinement or excluded oracle-failure regions can leave a midpoint-dependent support factor; finite normalization alone does not remove it. Smooth confinement included in V on the full space is distinct from a hard domain restriction.

Remark on the choice of coupling. With the rotation-minimized d^2 in place of d_{c}^2, Π would be invariant to a separate rotation of either walker, and the relative orientation would carry no weight. The midpoint would then have to be formed after rotating one walker onto the other. This map is not linear, and equation (S.23) would hold only approximately. The common frame, together with move (b), restores the exact identity at the cost of one oracle-free rotation per sweep. The output frame of a pair, in which both walkers are turned by one common rotation before they are written, leaves Π unchanged by equation (S.14).

3. PROOF OF STATIONARITY AND OF THE DIFFUSION LIMIT

The following proofs separate detailed balance of the fixed-target kernel from its small-step approximation to overdamped dynamics. The latter motivates equation (10) but does not establish equality of finite-step and diffusion committors.

3.1. Detailed balance of delayed acceptance

We prove that the move satisfies detailed balance with respect to π of equation (4) for any fixed c, b and $\tilde{V}$. Let $g(\mathbf{y} \mid \mathbf{x}) = g(\mathbf{x} \mid \mathbf{y})$ be a symmetric proposal and write $\alpha = \prod_{k=1}^{3}\min(1, r_k)$ with

$$r_1 = e^{\Delta c}, \qquad r_2 = e^{-\beta\Delta(\tilde{V}+b)}, \qquad r_3 = e^{-\beta\Delta(V-\tilde{V})}, \tag{S.24}$$

where $\Delta f = f(\mathbf{y}) - f(\mathbf{x})$. The reverse move has ratios $1/r_k$, and $\min(1, r)/\min(1, 1/r) = r$ for every $r > 0$. Hence

$$\frac{\alpha(\mathbf{x}, \mathbf{y})}{\alpha(\mathbf{y}, \mathbf{x})} = r_1 r_2 r_3 = \exp\big[\Delta c - \beta\Delta V - \beta\Delta b\big] = \frac{\pi(\mathbf{y})}{\pi(\mathbf{x})}, \tag{S.25}$$

because $\tilde{V}$ cancels between r_2 and r_3. With g symmetric this gives, for all $\mathbf{x} \neq \mathbf{y}$,

$$\pi(\mathbf{x})\, g(\mathbf{y} \mid \mathbf{x})\, \alpha(\mathbf{x}, \mathbf{y}) = \pi(\mathbf{y})\, g(\mathbf{x} \mid \mathbf{y})\, \alpha(\mathbf{y}, \mathbf{x}), \tag{S.26}$$

which is detailed balance[41]. Each move of a fixed sequence preserves π, and so does their composition[63]. A failed oracle call sets $V(\mathbf{y}) = +\infty$ and $r_3 = 0$. If the oracle fails on a fixed set $\mathcal{F}$, equation (S.26) then holds for the target $\pi\!\upharpoonright_{\mathcal{F}^{\mathrm{c}}}$, and the chain samples π restricted to the complement of $\mathcal{F}$. The order of the three tests does not enter the proof and is chosen so that the cheapest test comes first. If the oracle fails at random rather than on a fixed set, this restriction is not defined, and the retries of Supplementary Section 5.1 serve to make such failures rare.

3.2. Single versus repeated pre-screen

We prove that redrawing trials until one passes the pre-screen changes the stationary distribution. Consider the update of walker n at fixed $\mathbf{R}_m$, with conditional target $\pi_n(\mathbf{x}) = Z_n^{-1}\exp[c(\mathbf{x}) - \beta_{\mathrm{b}}V(\mathbf{x})]$, where $c(\mathbf{x}) = -d_{\mathrm{c}}^2(\mathbf{x}, \mathbf{R}_m)/(\hbar\tau_{12})$ and Z_n normalizes π_n. A trial passes the pre-screen with probability $a(\mathbf{x}, \mathbf{y}) = \min(1, e^{c(\mathbf{y})-c(\mathbf{x})})$, and one draw passes with probability

$$A(\mathbf{x}) = \int g(\mathbf{y} \mid \mathbf{x})\, a(\mathbf{x}, \mathbf{y})\, \mathrm{d}\mathbf{y}. \tag{S.27}$$

Suppose trials are redrawn until one passes, for at most K draws, and a passing trial is then accepted with $m(\mathbf{x}, \mathbf{y}) = \min(1, e^{-\beta_{\mathrm{b}}\Delta V})$. A trial $\mathbf{y} \neq \mathbf{x}$ is accepted at the k-th draw with probability $[1 - A(\mathbf{x})]^{k-1} g(\mathbf{y} \mid \mathbf{x}) a(\mathbf{x}, \mathbf{y}) m(\mathbf{x}, \mathbf{y})$. Summing over $k \leq K$ gives the transition density

$$T(\mathbf{x}, \mathbf{y}) = g(\mathbf{y} \mid \mathbf{x})\, a(\mathbf{x}, \mathbf{y})\, m(\mathbf{x}, \mathbf{y})\, \frac{1 - [1 - A(\mathbf{x})]^K}{A(\mathbf{x})}. \tag{S.28}$$

Now define

$$\pi_n^{(K)}(\mathbf{x}) = \frac{1}{Z_n^{(K)}}\,\pi_n(\mathbf{x})\,\frac{A(\mathbf{x})}{1-[1-A(\mathbf{x})]^K}, \tag{S.29}$$

where $Z_n^{(K)}$ normalizes $\pi_n^{(K)}$. Then

$$Z_n Z_n^{(K)}\,\pi_n^{(K)}(\mathbf{x})\,T(\mathbf{x},\mathbf{y}) = g(\mathbf{y}\mid\mathbf{x})\min\!\big(e^{c(\mathbf{x})},e^{c(\mathbf{y})}\big) \times \min\!\big(e^{-\beta_{\mathrm{b}}V(\mathbf{x})},e^{-\beta_{\mathrm{b}}V(\mathbf{y})}\big), \tag{S.30}$$

which is symmetric in $\mathbf{x}$ and $\mathbf{y}$. The update therefore satisfies detailed balance with respect to $\pi_n^{(K)}$. For $K=1$, the single test of tReax_expedition, $\pi_n^{(1)}=\pi_n$. For $K\to\infty$, $\pi_n^{(\infty)}$ is π_n tilted by the pass probability A. A is not constant. Near the other walker almost every displacement increases d_{c}^2, whereas far from it about half of the displacements decrease it. MCTSS uses the repeated pre-screen as a search device[6], for which such a tilt is harmless. The midpoint weights and the transition-state ensemble require Π exactly, which the single test provides at the price of about $1/A$ more steps.

3.3. Kinetic factor in the acceptance

We prove that placing the kinetic factor in the acceptance breaks detailed balance with respect to π_τ of equation (2). The rule obtained from equation (S.1) accepts with

$$\alpha_{\mathrm{M}}(\mathbf{x},\mathbf{y}) = \min\!\big\{1,\exp\!\big[-\beta\Delta V - k(\mathbf{x},\mathbf{y})\big]\big\}, \tag{S.31}$$

where $k = d^2(\mathbf{x},\mathbf{y})/(\hbar\tau) > 0$ for $\mathbf{x}\neq\mathbf{y}$. Let Z be the normalization of π_τ. Then

$$\begin{aligned} Z\,\pi_\tau(\mathbf{x})\,\alpha_{\mathrm{M}}(\mathbf{x},\mathbf{y}) &= \min\!\big\{e^{-\beta V(\mathbf{x})},\,e^{-\beta V(\mathbf{y})-k}\big\},\\ Z\,\pi_\tau(\mathbf{y})\,\alpha_{\mathrm{M}}(\mathbf{y},\mathbf{x}) &= \min\!\big\{e^{-\beta V(\mathbf{y})},\,e^{-\beta V(\mathbf{x})-k}\big\}. \end{aligned} \tag{S.32}$$

For any pair with $\beta V(\mathbf{x}) \geq \beta V(\mathbf{y})+k$, the first line equals $e^{-\beta V(\mathbf{x})}$ and the second $e^{-\beta V(\mathbf{x})-k}$, which is smaller. Multiplying both by the symmetric proposal density shows that the flux from $\mathbf{x}$ to $\mathbf{y}$ exceeds the reverse flux. Removing the kinetic factor from the acceptance gives $\alpha=\min(1,e^{-\beta\Delta V})$, a special case of Supplementary Section 3.1. The kinetic factor then acts once, through the proposal (Supplementary Section 1.2).

3.4. Small-step limit of the move

We prove that, for small displacements, the move generates overdamped Langevin dynamics. Let $U=V+b$ be twice continuously differentiable, let f have bounded continuous derivatives up to third order, and let P be the kernel of the single-walker move at parameter τ without coupling or surrogate. Fix atom i and write $\mathbf{a}=\nabla_i f(\mathbf{R})$ and $\mathbf{u}=\nabla_i U(\mathbf{R})$, where ∇_i acts on the coordinates of atom i. Then $\Delta U = \mathbf{u}\cdot\boldsymbol{\delta} + O(\|\boldsymbol{\delta}\|^2)$, and $\min(1,e^{-x}) = 1 - x_+ + O(x^2)$ with $x_+=\max(x,0)$ gives $\alpha = 1-\beta(\mathbf{u}\cdot\boldsymbol{\delta})_+ + O(\|\boldsymbol{\delta}\|^2)$. Expanding f to second order,

$$\begin{aligned} \mathbb{E}\big[\alpha\,\Delta f\big] &= \mathbb{E}[\mathbf{a}\cdot\boldsymbol{\delta}] - \beta\,\mathbb{E}\big[(\mathbf{a}\cdot\boldsymbol{\delta})(\mathbf{u}\cdot\boldsymbol{\delta})_+\big]\\ &\quad + \tfrac{1}{2}\,\mathbb{E}\big[\boldsymbol{\delta}^{\mathsf{T}}\nabla_i\nabla_i^{\mathsf{T}} f\,\boldsymbol{\delta}\big] + O(\sigma_i^3). \end{aligned} \tag{S.33}$$

The first term vanishes. For the second, $(\mathbf{u}\cdot\boldsymbol{\delta})_+ = \frac{1}{2}(\mathbf{u}\cdot\boldsymbol{\delta}) + \frac{1}{2}|\mathbf{u}\cdot\boldsymbol{\delta}|$, and $\mathbb{E}[(\mathbf{a}\cdot\boldsymbol{\delta})|\mathbf{u}\cdot\boldsymbol{\delta}|]=0$ because the integrand is odd under $\boldsymbol{\delta}\to-\boldsymbol{\delta}$. Hence $\mathbb{E}[(\mathbf{a}\cdot\boldsymbol{\delta})(\mathbf{u}\cdot\boldsymbol{\delta})_+] = \frac{1}{2}\sigma_i^2\,\mathbf{a}\cdot\mathbf{u}$. The third term equals $\frac{1}{2}\sigma_i^2\nabla_i^2 f$. Averaging over the atom with probability $1/N$ gives

$$(Pf-f)(\mathbf{R}) = \sum_{i=1}^{N} D_i\big[\nabla_i^2 f - \beta\,\nabla_i U\cdot\nabla_i f\big](\mathbf{R}) + O(\sigma_{\max}^3), \tag{S.34}$$

where $\sigma_{\max}=\max_i\sigma_i$ and

$$D_i = \frac{\sigma_i^2}{2N} = \frac{\hbar\tau}{4Nm_i} \tag{S.35}$$

is the diffusion constant per step. The right-hand side of equation (S.34) is the generator of overdamped Langevin dynamics with mass-dependent diffusion constants[45, 64]. For invariant f, $f(A_{\mathbf{R}}(\mathbf{y})) = f(\mathbf{y})$, so the implemented kernel P_{A} has the same expansion. The error terms are small when $\beta\sigma_i\|\nabla_i U\| \ll 1$ and $\beta\sigma_i^2\|\nabla_i\nabla_i^{\mathsf{T}}U\| \ll 1$. At $\tau^*=25.6$ fs a hydrogen atom has $\sigma = 0.28$ Å, so the correspondence is approximate for molecules.

3.5. Variational form of the committor

We prove that the committor of the limiting diffusion minimizes K of equation (10). Let

$$\begin{aligned} \mathcal{L}f &= \sum_i D_i\big(\nabla_i^2 f - \beta^*\nabla_i V\cdot\nabla_i f\big)\\ &= e^{\beta^* V}\sum_i D_i\nabla_i\cdot\big(e^{-\beta^* V}\nabla_i f\big), \end{aligned} \tag{S.36}$$

with $D_i = \hbar\tau^*/(4Nm_i)$, and let Ω be the complement of the cores C_a and C_b, with no flux through its outer boundary. The committor q of this diffusion solves $\mathcal{L}q = 0$ in Ω with $q=0$ on ∂C_a and $q=1$ on ∂C_b[3]. For admissible f, which take these boundary values and are constant on each core,

$$I[f] = \int_\Omega e^{-\beta^* V}\sum_i D_i\|\nabla_i f\|^2\,\mathrm{d}\mathbf{R} = \frac{\hbar\tau^*}{4N}\,Z^*K[f], \tag{S.37}$$

where $Z^* = \int e^{-\beta^* V}\mathrm{d}\mathbf{R}$, since $\sum_i D_i\|\nabla_i f\|^2 = \frac{\hbar\tau^*}{4N}|\nabla_u f|^2$. Let h vanish on ∂C_a and ∂C_b. Integration by

parts, with no boundary terms, gives

$$\int_\Omega e^{-\beta^* V} \sum_i D_i \nabla_i q \cdot \nabla_i h \, \mathrm{d}\mathbf{R} = -\int_\Omega h \, e^{-\beta^* V} \mathcal{L} q \, \mathrm{d}\mathbf{R} = 0. \tag{S.38}$$

Hence $I[q+h] = I[q] + I[h] \geq I[q]$. Therefore, q minimizes I and hence K, and the prefactor $\hbar\tau^*/(4N)$ does not affect the minimizer[19, 65]. The relation between the Monte Carlo committor, equation (9), and $K[q]$ therefore relies on the small-step limit of Supplementary Section 3.4. The diffusion functional approximates the finite-step Monte Carlo target. Shooting provides a separate diagnostic under the specified kernel and cores, subject to the implementation and finite-budget limitations in Supplementary Section 8.4.

4. CONDITIONAL CONTINUOUS-PATH ENERGY BOUND

Equation (8) records evaluated energies. A continuous-path bound additionally requires control of the unevaluated segments; the derivation below identifies this correction.

4.1. Segments

The minimax energy between two configurations is

$$E^*(\mathbf{x}, \mathbf{y}) = \inf_\gamma \max_{s\in[0,1]} V\big(\gamma(s)\big), \tag{S.39}$$

where the infimum runs over continuous paths γ with $\gamma(0) = \mathbf{x}$ and $\gamma(1) = \mathbf{y}$. V is called L_V-Lipschitz on a set $\mathcal{S}$ when $|V(\mathbf{x}) - V(\mathbf{y})| \leq L_V \|\mathbf{x} - \mathbf{y}\|_{\mathsf{W}}$ for all $\mathbf{x}$ and $\mathbf{y}$ in $\mathcal{S}$.

Straight segment. Let V be L_V-Lipschitz on the segment from $\mathbf{x}$ to $\mathbf{y}$, of length $\ell_{xy} = \|\mathbf{y} - \mathbf{x}\|_{\mathsf{W}}$. At parameter s, $V \leq V(\mathbf{x}) + L_V s \ell_{xy}$ and $V \leq V(\mathbf{y}) + L_V(1-s)\ell_{xy}$. The smaller of the two bounds is at most their mean, so

$$\begin{aligned}\max_{s\in[0,1]} V\big(\mathbf{x} + s(\mathbf{y}-\mathbf{x})\big) &\leq \tfrac{1}{2}\big[V(\mathbf{x}) + V(\mathbf{y}) + L_V \ell_{xy}\big] \\ &\leq \max\{V(\mathbf{x}), V(\mathbf{y})\} + \tfrac{1}{2} L_V \ell_{xy}.\end{aligned} \tag{S.40}$$

Rigid motion. Let $\mathbf{y}' = g\mathbf{y}$ with $g = (\mathsf{Q}, \mathbf{t})$. $SO(3)$ is path-connected, since $\mathsf{Q} = \exp(\mathbf{\Omega})$ for a skew-symmetric $\mathbf{\Omega}$. The path $s \mapsto g(s)\mathbf{y}$ with $g(s) = (\exp(s\mathbf{\Omega}), s\mathbf{t})$ joins $\mathbf{y}$ to $\mathbf{y}'$, and by equation (S.6)

$$V\big(g(s)\mathbf{y}\big) = V(\mathbf{y}) \qquad \text{for all } s \in [0,1]. \tag{S.41}$$

4.2. Chains

We prove that, for a certified chain with bound E_{up} and end configurations $\mathbf{x}_a$ and $\mathbf{x}_b$ of its two end quenches,

$$E^*(\mathbf{x}_a, \mathbf{x}_b) \leq E_{\mathrm{up}} + \tfrac{1}{2} L_V h. \tag{S.42}$$

Here h is the largest mass-weighted length of a single accepted move, taken before frame removal, and of the spacing between consecutive evaluated points of a link, over all quenches and links that enter equation (8). V is assumed to be L_V-Lipschitz on a set that contains all these segments.

Construction of the path. The path Γ traverses the end quench of S_a backwards from $\mathbf{x}_a$ to the end node. It then follows the parent tree of the main nodes to the seed, continues to the end node on the other side and follows the end quench of S_b to $\mathbf{x}_b$. Each segment between a node and its parent is traversed along its link or along its quench trajectory, whichever has the lower maximum, as in equation (8). Let $\mathbf{y}_0 = \mathbf{x}_a, \mathbf{y}_1, \ldots, \mathbf{y}_M = \mathbf{x}_b$ be the evaluated configurations in the order in which Γ visits them. Two consecutive configurations are joined in one of two ways. For an accepted single-atom move, $\mathbf{y}_{k+1} = A_{\mathbf{y}_k}(\mathbf{y}_k + \mathbf{e}_i \otimes \boldsymbol{\delta})$, and Γ follows the straight segment to $\mathbf{y}_k + \mathbf{e}_i \otimes \boldsymbol{\delta}$, of length at most h, and then the rigid path of equation (S.41). Along a link, Γ follows the straight segment between consecutive evaluated points, again of length at most h. A link written in the frame of its first node is joined to the next node by a rigid path.

Induction along the path. Let Γ_k be the part of Γ from $\mathbf{y}_0$ to $\mathbf{y}_k$. We show by induction that

$$\max_{\Gamma_k} V \leq \max_{j\leq k} V(\mathbf{y}_j) + \tfrac{1}{2} L_V h. \tag{S.43}$$

- *Base case ($k = 0$):* Γ_0 is the single point $\mathbf{y}_0$, and $\max_{\Gamma_0} V = V(\mathbf{y}_0)$.
- *Induction hypothesis:* Equation (S.43) holds for k.
- *Induction step:* Γ_{k+1} is Γ_k followed by the piece from $\mathbf{y}_k$ to $\mathbf{y}_{k+1}$. On the straight part of this piece, equation (S.40) bounds V by $\max\{V(\mathbf{y}_k), V(\mathbf{y}_{k+1})\} + \frac{1}{2} L_V h$, since the displaced trial and $\mathbf{y}_{k+1}$ have the same energy. On the rigid part, V is constant by equation (S.41). Together with the induction hypothesis,

$$\max_{\Gamma_{k+1}} V \leq \max_{j\leq k+1} V(\mathbf{y}_j) + \tfrac{1}{2} L_V h. \tag{S.44}$$

By induction, equation (S.43) holds for $k = M$. The largest evaluated energy along Γ is the seed energy, the maximum over segments of the chosen link or trajectory maximum, or an end-descent maximum, which is E_{up} by equation (8). Restraints change which moves are accepted but not the energies recorded along Γ. Therefore, $\max_\Gamma V \leq E_{\mathrm{up}} + \frac{1}{2} L_V h$, and equation (S.42) follows from equation (S.39).

The code does not compute L_V or h. Under the stated assumptions, $\frac{1}{2} L_V h$ is the discretization correction. Smaller proposal widths and link spacings can reduce h, but their nominal values do not provide a deterministic bound on the largest accepted Gaussian move.

The evaluated straight-link spacing is at most $\ell/4$. The bound refers to $\mathbf{x}_a$ and $\mathbf{x}_b$. The end quench may reach a minimum of S_a other than the lowest one recorded for that state, so barriers are quoted against the lowest minimum found at the end of the search.

4.3. Routes

We finally extend the bound to routes through intermediate states. Concatenating a path from $\mathbf{x}$ to $\mathbf{y}$ with one from $\mathbf{y}$ to $\mathbf{z}$ gives a path from $\mathbf{x}$ to $\mathbf{z}$ whose maximum is the larger of the two maxima, so that

$$E^*(\mathbf{x},\mathbf{z}) \le \max\{E^*(\mathbf{x},\mathbf{y}), E^*(\mathbf{y},\mathbf{z})\}. \tag{S.45}$$

Let a route $A = S_0, S_1, \dots, S_K = B$ use certified chains with bounds $E_{\mathrm{up},k}$ and spacings h_k, and let $U = \max_k(E_{\mathrm{up},k} + \frac{1}{2}L_V h_k)$. Suppose that, within every intermediate state S_k, the end configurations of the two adjacent chains satisfy $E^* \le U$. Applying equation (S.45) along the route, with equation (S.42) for every chain, shows that the end configurations of the route in A and B satisfy $E^* \le U$. A route that minimizes the largest edge value on the graph of states is found by a variant of Dijkstra's algorithm for maximum-capacity paths[43]. The connection assumption within intermediate states is not checked by the code. It can fail when one bond graph contains several conformers separated by a barrier higher than U.

5. QUANTUM-INSPIRED MONTE CARLO SIMULATION

This section supplies the implementation details for Methods Sections 4.1–4.4. Supplementary Sections 5.1 and 5.2 specify the single-walker kernel and coupled-pair ensemble used in Methods Sections 4.1 and 4.2; their statistical properties are established in Supplementary Sections 2 and 3. Supplementary Sections 5.4–5.7 give the quench, state-assignment and chain-construction details for Methods Section 4.3. Supplementary Tables S.2 and S.3 collect the corresponding settings.

5.1. Proposal and delayed acceptance

This section describes how the move of Methods Section 4.1 is implemented. Its exactness is proved in Supplementary Section 3 and its invariance in Supplementary Section 2.

Target and tests. The target of one walker is equation (4). A symmetric proposal $g(\mathbf{y} \mid \mathbf{x}) = g(\mathbf{x} \mid \mathbf{y})$ generates the trial $\mathbf{y}$, which is accepted with

$$\alpha(\mathbf{x},\mathbf{y}) = \prod_{k=1}^{3} \min(1, r_k), \tag{S.46}$$

where $r_1 = e^{\Delta c}$, $r_2 = e^{-\beta\Delta(\tilde{V}+b)}$, $r_3 = e^{-\beta\Delta(V-\tilde{V})}$ and $\Delta f = f(\mathbf{y}) - f(\mathbf{x})$. Each test draws its own uniform number, and the first rejection ends the move. Absent terms are zero, and without a surrogate the second and third tests merge into $\min(1, e^{-\beta\Delta(V+b)})$. Supplementary Section 3.1 shows that π is stationary whatever the surrogate, as long as c, b and $\tilde{V}$ are fixed while the chain runs.

Failures. A failed oracle call is treated as $V(\mathbf{y}) = +\infty$, so $r_3 = 0$ and the trial is rejected. A failed surrogate evaluation returns $+\infty$ and is rejected by the second test without an oracle call. A chain is not started from a configuration where the surrogate is infinite, since it could never leave it. During coupled-walker exploration the oracle is called up to four times before a failure is declared, which absorbs transient failures of file-driven programs.

Single walker. One atom is chosen with probability $1/N$ and displaced by $\boldsymbol{\delta} \sim \mathcal{N}(\mathbf{0}, \sigma_i^2 \mathrm{I}_3)$. For free molecules the trial is then centred at its centre of mass and rotated onto the current structure by the proper mass-weighted Kabsch rotation[37, 38]. By Supplementary Section 2.4 this frame removal leaves the distribution of every invariant quantity unchanged. Frame removal is skipped for periodic cells and for systems with fixed atoms.

Coupled walkers. For a pair, walker n is updated with the conditional target

$$\pi(\mathbf{R}_n \mid \mathbf{R}_m) = \frac{1}{Z_n} \exp\Big[-\beta_{\mathrm{b}} V(\mathbf{R}_n) - \frac{d_{\mathrm{c}}^2(\mathbf{R}_n, \mathbf{R}_m)}{\hbar\tau_{12}}\Big], \tag{S.47}$$

where Z_n is the integral of the exponential over $\mathbf{R}_n$ at fixed $\mathbf{R}_m$ and d_{c}^2 is the common-frame distance of equation (5). The coupling enters as c. Trials of either walker are centred but never rotated, so both walkers stay in one frame. After both walkers have been updated, the second walker receives the rigid rotation of Supplementary Section 2.5, whose angle has standard deviation

$$\sigma_{\mathrm{rot}} = \min\Big\{0.5, \Big(\frac{\hbar\tau_{12}}{\bar{I}}\Big)^{1/2}\Big\} \text{ rad}, \tag{S.48}$$

where $\bar{I} = \frac{2}{3}\sum_i m_i \|\bar{\mathbf{r}}_{2,i}\|^2$ is the mean principal moment of inertia of the centred second walker. A rotation by this angle changes $d_{\mathrm{c}}^2/(\hbar\tau_{12})$ by about one. The rotation keeps the energy, so it is tested on the coupling alone and needs no oracle call. The pre-screen is tested once, and a failed test leaves the walker in place. MCTSS instead repeats the pre-selection until a trial passes, for up to 100 attempts[6], which tilts the stationary distribution (Supplementary Section 3.2). The squared distance is carried between updates, because it changes only when a move is accepted.

Adaptive widths. Quenches and the tuning phase of transition-state ensemble sampling adapt the proposal width. Such chains do not satisfy detailed balance and are used as optimizers or discarded. For ensemble sampling, the width is frozen before recording. The resulting

fixed kernel preserves the target; equilibrium estimates additionally require equilibration and adequate mixing.

5.2. Pair ensemble and midpoint weights

Let $\bar{\mathbf{R}}_1$ and $\bar{\mathbf{R}}_2$ be the centred walkers in their common frame and write $\mathbf{M} = (\bar{\mathbf{R}}_1 + \bar{\mathbf{R}}_2)/2$ and $\mathbf{D} = \bar{\mathbf{R}}_1 - \bar{\mathbf{R}}_2$, a linear change of variables with a Jacobian of unit modulus. Then $d_c^2 = \mathbf{D}^\mathsf{T}\mathsf{W}\mathbf{D}$, where W is the diagonal mass matrix. Multiplying Π of equation (5) by the weight w of equation (6) gives equation (S.23), in which the two factors separate. After normalization, the weighted marginal of $\mathbf{M}$ is exactly

$$p_w(\mathbf{M}) = \frac{\exp\bigl[-2\beta_\mathrm{b}V(\mathbf{M})\bigr]}{\displaystyle\int \exp\bigl[-2\beta_\mathrm{b}V(\mathbf{M}')\bigr]\,\mathrm{d}\mathbf{M}'}, \tag{S.49}$$

the ensemble of $\tau_\mathrm{mid} = 2\tau_\mathrm{b}$. The identity holds for free molecules as well, because the midpoint is formed in the common frame without rotating one walker onto the other (Supplementary Section 2.5). Recorded midpoints are rotated onto the first one before they are written, which changes only the frame of the output.

Unweighted marginal. Without weights, expanding to second order in $\mathbf{D}$ gives $V(\mathbf{R}_1) + V(\mathbf{R}_2) = 2V(\mathbf{M}) + \frac{1}{4}\mathbf{D}^\mathsf{T}\mathsf{H}\mathbf{D} + O(\|\mathbf{D}\|^4)$, where H is the Hessian at $\mathbf{M}$. The Gaussian integral over $\mathbf{D}$ in d dimensions yields

$$\begin{aligned} p(\mathbf{M}) \approx{}& \frac{\pi^{d/2}}{Z_\Pi}\, e^{-2\beta_\mathrm{b}V(\mathbf{M})} \\ &\times \det\Bigl[\frac{\mathsf{W}}{\hbar\tau_{12}} + \frac{\beta_\mathrm{b}}{4}\mathsf{H}(\mathbf{M})\Bigr]^{-1/2}. \end{aligned} \tag{S.50}$$

The determinant is the factor that the weights remove.

Distribution of the log-weights. To leading order in τ_{12}, $\mathbf{D}$ is Gaussian with covariance $\mathbf{\Sigma}_D = \frac{1}{2}\hbar\tau_{12}\mathsf{W}^{-1}$, and $\ln w \approx \frac{1}{4}\beta_\mathrm{b}\mathbf{D}^\mathsf{T}\mathsf{H}\mathbf{D}$. The moments of a Gaussian quadratic form then give

$$\begin{aligned} \langle \ln w\rangle &\approx \tfrac{1}{4}\beta_\mathrm{b}\,\mathrm{tr}(\mathsf{H}\mathbf{\Sigma}_D) = \frac{n_{12}\tau_\mathrm{b}^2}{8}\sum_k \omega_k^2, \\ \mathrm{Var}(\ln w) &\approx \tfrac{1}{8}\beta_\mathrm{b}^2\,\mathrm{tr}\bigl[(\mathsf{H}\mathbf{\Sigma}_D)^2\bigr] = \frac{(n_{12}\tau_\mathrm{b}^2)^2}{32}\sum_k \omega_k^4, \end{aligned} \tag{S.51}$$

where ω_k^2 are the eigenvalues of $\mathsf{W}^{-1/2}\mathsf{H}\mathsf{W}^{-1/2}$ and $\beta_\mathrm{b}\hbar\tau_{12} = n_{12}\tau_\mathrm{b}^2$ was used. For a Gaussian log-weight distribution, the large-sample Kish fraction tends to $\exp[-\mathrm{Var}(\ln w)]$. The finite-sample quantity is defined in Supplementary Section 5.9 and does not account for autocorrelation. The weights therefore stay useful when $n_{12}(\omega_k\tau_\mathrm{b})^2$ is of order one or smaller for the stiffest modes. For a C–H stretch near 3000 cm^{-1} at $\tau_\mathrm{b} = \tau^*/2 = 12.8$ fs, $n_{12}(\omega\tau_\mathrm{b})^2 \approx 2.6$ for $n_{12} = 0.05$, compared with about 160 for the exploration value $n_{12} = 3$. This motivates the stiff coupling used for the transition-state ensemble. It also shows that the weights of exploration pairs, with $n_{12} = 3$, are dominated by few samples, so the weighted estimates of Supplementary Section 9.8 are reported only as heuristics.

Meeting criterion. If all atoms were displaced at once, the mean mass-weighted squared displacement per step would be $\frac{3}{2}N\hbar\tau$[6]. Following MCTSS, a molecular run stops when $d^2(\mathbf{R}_1,\mathbf{R}_2)/(\hbar\tau_\mathrm{b}) \le \frac{3}{2}N s_\mathrm{meet}$ with $s_\mathrm{meet} = 1$. For the two-dimensional models the test is $d^2/(\hbar\tau_\mathrm{b}) \le s_\mathrm{meet}$ with $s_\mathrm{meet} = 0.5$.

Overlap criterion. The progress of a structure is $p(\mathbf{R}) = \langle \mathbf{R} - \mathbf{R}_A^0, \mathbf{R}_B^0 - \mathbf{R}_A^0\rangle_\mathsf{W}/\|\mathbf{R}_B^0 - \mathbf{R}_A^0\|_\mathsf{W}^2$, evaluated after alignment onto an endpoint, with $\langle \mathbf{a},\mathbf{b}\rangle_\mathsf{W} = \mathbf{a}^\mathsf{T}\mathsf{W}\mathbf{b}$. It is 0 at the reactant and 1 at the product. A run also stops once the largest progress reached by walker 1 is at least the smallest progress reached by walker 2. For the two-dimensional models the axis is the unweighted line between the two starting points.

5.3. Coupled-walker exploration ladder

Supplementary Algorithm 1 gives the ladder. A rung runs four independent MCTSS runs of at most 5000 steps. Molecular runs use $n_{12} = 3$ and the two-dimensional runs $n_{12} = 1$. Every step of a molecular run updates both walkers and then rotates the second one as in Supplementary Section 5.1. A run joins when it meets or overlaps. Both walkers are written in one frame, obtained by the common rotation that takes their midpoint onto the starting midpoint, so consecutive frames of the joined path share one orientation. The exploration parameter τ_x is the largest τ_b of a successful rung before the first failure. Its only role downstream is to set proposal widths and the length ℓ. It is not used as a temperature in any later step. The paths labelled by quenching come from converged runs of all rungs with $\tau_\mathrm{b} \le \tau_\mathrm{x}$, coldest rung first and alternating between rungs, so that more than one effective temperature contributes.

5.4. Quenches

A quench runs the move of Methods Section 4.1 for one walker at each thermal energy of its schedule, starting from the hottest. Its initial proposal parameter is given in Supplementary Table S.3. After every block of 25 steps, the proposal parameter is halved if the block acceptance was below 0.2 and multiplied by 1.5 if it exceeded 0.5. When patience $P > 0$ is set, a level ends after P consecutive blocks in which the lowest value of $V + b$ fell by less than 0.1 meV. A quench is flagged as converged when its coldest level ended in this way. Its result is the configuration held at the end, together with the highest energy held along the way.

Single-atom moves relax one atom at a time, so budgets grow with the number of degrees of freedom d ($3N$ for molecules, 2 in two dimensions). With $r = \max(1, d/18)$, every schedule has r times the steps per

level listed in Supplementary Table S.3. The schedules marked as cooled also divide their thermal energies by r, so that the excess energy $\frac{d}{2}k_{\mathrm{B}}T$ of the hottest level stays that of an 18-dimensional system. The optional per-system multipliers of budget and patience were not changed from their defaults.

5.5. States and majority-quench labels

Fingerprints. For reactions, atoms i and j are bonded when $r_{ij} < r_i^{\mathrm{cov}} + r_j^{\mathrm{cov}} + 0.4$ Å, with the covalent radii of ASE[56]. An additive tolerance separates H–H distances of 0.74 to 0.80 Å and stretched from broken ring bonds, which a multiplicative factor cannot do for both. Two bond graphs belong to the same state only if they are identical. For alanine dipeptide the fingerprint is (ϕ, ψ), and two configurations share a state when both periodic differences are at most 60°. For the two-dimensional models two minima share a state when they lie within the basin radius of Supplementary Table S.5.

Registration. A quenched configuration that matches a known state is assigned to it, and the state keeps the lower of the two minima. An unmatched configuration registers a new intermediate I_k. Two exceptions apply. If the quench did not converge, or if new states are disallowed as for alanine dipeptide, the configuration is assigned to the nearest known state wherever the fingerprint defines a finite distance. Bond graphs have no such distance, so a non-converged quench with a new bond graph still registers a new state. A label with several quenches takes the majority, and ties go to the label met first.

Bisection. The two ends of an exploration path carry the labels A and B. For every interval whose end labels differ, the middle frame is quenched and both halves are examined again. The recursion stops at adjacent frames. The number of quenches grows logarithmically with the path length for each label change. Frames that lie between two quenched frames of the same state inherit that state and provide the training data of the label classifier.

5.6. Progress coordinates

Let $\mathcal{B}_a$ and $\mathcal{B}_b$ be the bond sets of the minima of S_a and S_b. The bond-change coordinate is

$$\xi_{\mathrm{bond}}(\mathbf{R}) = \sum_{(i,j)\in\mathcal{B}_a\setminus\mathcal{B}_b} r_{ij} - \sum_{(k,l)\in\mathcal{B}_b\setminus\mathcal{B}_a} r_{kl}, \qquad \text{(S.52)}$$

which generalizes the $\mathrm{S_N2}$ coordinate $r(\mathrm{C{-}Br}) - r(\mathrm{C{-}Cl})$. For a dihedral fingerprint $\boldsymbol{\theta}$, let $\boldsymbol{\Delta} = \mathrm{wrap}(\boldsymbol{\theta}_b - \boldsymbol{\theta}_a)$ with wrap mapping angles into $[-180°, 180°)$. Then

$$\xi_{\mathrm{dih}}(\mathbf{R}) = \frac{\mathrm{wrap}\big(\boldsymbol{\theta}(\mathbf{R}) - \boldsymbol{\theta}_a\big)\cdot\boldsymbol{\Delta}}{\boldsymbol{\Delta}\cdot\boldsymbol{\Delta}}, \qquad \text{(S.53)}$$

which is 0 at S_a and 1 at S_b. When neither coordinate exists, for instance for the two-dimensional models, ξ_{label} is the latent z of a classifier with the architecture of the committor network (Supplementary Section 8.1). It is trained on the path frames labelled S_a or S_b and on the two minima, with targets 0 and 1. The loss is the binary cross-entropy of $\sigma(3z)$ with the two classes weighted equally. Optimization uses 400 full-batch epochs of Adam[59] with learning rate 3×10^{-3} and weight decay 10^{-3}, which keeps z finite on separable data so that its level sets remain usable as slices. When chains are regrown after committor learning, the coordinate is the committor latent z_θ.

5.7. Continuation chains and the energy bound

This section completes the chain construction and the bound E_{up} of Methods Section 4.3, and Supplementary Algorithm 2 summarizes it. The seed is restrained at its own coordinate, clipped to the interval between the two minima. Slices farther than $\Delta\xi/4$ from the seed on either side are visited in order, each node starting from the previous one. For each node, the restraint centre moves from the current ξ to the target in five equal steps. Each step is a short quench at the second-coldest level of the node schedule with half its step budget. The full node schedule then follows, and the restraint is stiffened fourfold if the node misses its slice by more than $\Delta\xi/4$, up to three times. The proposal parameter carried to the next node is the larger of the value the quench ended with and a quarter of the value it started with. An inserted node targets the mid-point of its two neighbours with half the slice width and the current restraint. Densification stops after at most three recursive levels.

Candidates for the bound. The main nodes are the seed node and the nodes on the slice grid. Every main node m other than the seed has a parent p, the main node from which it was grown. The segment between m and p spans the node indices $a = \min(m,p)$ to $b = \max(m,p)$, including inserted nodes. Its straight-path maximum E^{link} is the largest node energy in $[a,b]$ and interior link energy between them. Links are evaluated in the frame of their first node. Its trajectory maximum E^{traj} is the largest energy held along all quenches that grew m from p, including the starting energy. The chain bound is equation (8), and the location of its maximum is reported as a node, a straight segment, a trajectory segment or an end descent.

What the bound certifies. Supplementary Section 4.2 shows that the minimax energy between the end configurations of a certified chain is at most $E_{\mathrm{up}} + \frac{1}{2}L_V h$, where L_V is a Lipschitz constant of V and h the largest spacing between consecutive evaluated configurations. Certification checks endpoint connectivity and the largest evaluated energy; the code does not compute L_V or the discretization correction. The reported barrier convention is defined in Methods Section 4.3; it uses the lowest

recorded state minimum even when a chain ends at another conformer.

Channels and routes. Two certified chains form one channel when the highest node of each lies within 3ℓ of some node of the other. The lowest chain of each channel is written out and used to locate the quench-label crossing and to sample the transition-state ensemble. On the undirected graph whose edges carry the lowest E_{up} of each pair, a Dijkstra-type search that minimizes the largest edge value along the path[43] returns the A-to-B route and its bound. The route value bounds the minimax energy between A and B under the connection assumption of Supplementary Section 4.3.

5.8. Quench-label crossing and curvature

Crossing labels are majority votes of three quenches with the label schedule. Bisection first narrows the pair of adjacent nodes whose labels differ, starting from the known end labels. The straight link between them is then bisected four times, leaving a bracket with one sixteenth of the original link length. The candidate is its midpoint, and 1/32 denotes half this geometric width rather than uncertainty in a basin boundary. Because the crossing lies on a straight link, $E_0^{\ddagger}$ is not bounded by E_{up}, which may be set by a trajectory segment. The energy curvature along the chain is

$$k = \frac{V(\mathbf{R}_k + h\hat{\mathbf{e}}) + V(\mathbf{R}_k - h\hat{\mathbf{e}}) - 2V(\mathbf{R}_k)}{h^2}, \qquad \text{(S.54)}$$

where $\mathbf{R}_k$ is the highest node, $\hat{\mathbf{e}}$ the unit mass-weighted tangent through its neighbours and $h = 0.05$ amu$^{1/2}$ Å. The implied frequency is $\omega = \sqrt{|k|}$, reported with the sign of k as a wavenumber, and $T_{\mathrm{c}} = \hbar|\omega|/(2\pi k_{\mathrm{B}})$. The tangent is not a normal mode, so both values are diagnostics.

The geometric bisection interval does not quantify stochastic label uncertainty. With three random quench votes per query, inconsistent labels can alter the bracket; its midpoint is a protocol-defined candidate rather than verified membership of a unique separatrix.

5.9. Local straddling-pair ensemble (TSE)

Let $s(\mathbf{R}) = \pm[\xi(\mathbf{R}) - \xi^*]$, where $\xi^* = \xi(\mathbf{R}_0^{\ddagger})$ and the sign is chosen so that $s < 0$ at the minimum of S_a. The ensemble is sampled from Π of equation (5) with $\tau_{\mathrm{b}} = \tau^*/2$ and $\tau_{12} = 0.05\tau_{\mathrm{b}}$, multiplied by the indicator $\mathbb{1}[s(\mathbf{R}_1) < 0 < s(\mathbf{R}_2)]$. The indicator enters the geometric test of the move, so a move that breaks it costs no oracle call. The walkers share one frame and their relative orientation is sampled by the rotation of Supplementary Section 5.1. Because s is invariant (Supplementary Section 2.3), the rotation never violates the indicator. During 2000 tuning sweeps the proposal parameter starts at τ_{12} and is multiplied by $\exp[2(a - 0.3)]$ after every 100 moves with acceptance a. The width is then frozen and 20000 sweeps are recorded. The midpoint energy is evaluated every 10 recorded sweeps.

Weighted joint density. Since $2\beta_{\mathrm{b}} = \beta^*$, equation (S.23) gives the weighted joint density

$$\Pi\, w\, \mathbb{1} = \frac{1}{Z_\Pi}\, e^{-\beta^* V(\mathbf{M})} \exp\Big[-\frac{\mathbf{D}^{\mathsf{T}}\mathsf{W}\mathbf{D}}{\hbar\tau_{12}}\Big] \times \mathbb{1}\big[s(\mathbf{M} + \tfrac{1}{2}\mathbf{D}) < 0 < s(\mathbf{M} - \tfrac{1}{2}\mathbf{D})\big], \qquad \text{(S.55)}$$

where $\mathbf{M}$ and $\mathbf{D}$ are defined in Supplementary Section 5.2. The marginal derivation below additionally assumes the full-space relative-coordinate integration used in Supplementary Section 2.5; within that setting, it linearizes s over the scale of $\mathbf{D}$.

Marginal of the midpoint. At fixed $\mathbf{M}$, the second factor of equation (S.55) makes $\mathbf{D}$ Gaussian with covariance $\mathbf{\Sigma}_D = \frac{1}{2}\hbar\tau_{12}\mathsf{W}^{-1}$. With $\mathbf{g} = \nabla s(\mathbf{M})$ and $s(\mathbf{M} \pm \frac{1}{2}\mathbf{D}) \approx s(\mathbf{M}) \pm \frac{1}{2}\mathbf{g}\cdot\mathbf{D}$, the indicator requires $-\mathbf{g}\cdot\mathbf{D} > 2|s(\mathbf{M})|$. The scalar $\mathbf{g}\cdot\mathbf{D}$ is Gaussian with mean zero and variance $\sigma_g^2 = \mathbf{g}^{\mathsf{T}}\mathbf{\Sigma}_D\mathbf{g} = \frac{1}{2}\hbar\tau_{12}\,\mathbf{g}^{\mathsf{T}}\mathsf{W}^{-1}\mathbf{g}$, so the indicator holds with probability $\frac{1}{2}\operatorname{erfc}(\sqrt{2}|s(\mathbf{M})|/\sigma_g)$. The marginal of $\mathbf{M}$ is therefore

$$p(\mathbf{M}) = \frac{1}{Z_s}\, e^{-\beta^* V(\mathbf{M})}\, \tfrac{1}{2}\operatorname{erfc}\Big(\frac{\sqrt{2}\,|s(\mathbf{M})|}{\sigma_g(\mathbf{M})}\Big), \qquad \text{(S.56)}$$

where Z_s is the integral of the numerator over $\mathbf{M}$. The erfc factor is a window of width of order σ_g around the dividing surface.

Stiff limit. With $\int_0^\infty \operatorname{erfc}(u)\,\mathrm{d}u = 1/\sqrt{\pi}$,

$$\int_{-\infty}^{\infty} \tfrac{1}{2}\operatorname{erfc}\Big(\frac{\sqrt{2}\,|s|}{\sigma_g}\Big)\mathrm{d}s = \frac{\sigma_g}{\sqrt{2}}\int_0^\infty \operatorname{erfc}(u)\,\mathrm{d}u = \frac{\sigma_g}{\sqrt{2\pi}}. \qquad \text{(S.57)}$$

As $\tau_{12} \to 0$ the window therefore tends to $(\sigma_g/\sqrt{2\pi})\,\delta(s)$, and

$$p(\mathbf{M}) \to \frac{\sigma_g(\mathbf{M})\, e^{-\beta^* V(\mathbf{M})}\, \delta\big(\xi(\mathbf{M}) - \xi^*\big)}{\displaystyle\int \sigma_g(\mathbf{M}')\, e^{-\beta^* V(\mathbf{M}')}\, \delta\big(\xi(\mathbf{M}') - \xi^*\big)\,\mathrm{d}\mathbf{M}'}, \qquad \text{(S.58)}$$

where $\sigma_g = (\frac{1}{2}\hbar\tau_{12})^{1/2}|\nabla\xi|_{\mathsf{W}^{-1}}$ with $|\nabla\xi|_{\mathsf{W}^{-1}} = (\mathbf{g}^{\mathsf{T}}\mathsf{W}^{-1}\mathbf{g})^{1/2}$. In this stiff limit the midpoints sample the Boltzmann distribution on the dividing surface with a surface element weighted by $|\nabla\xi|_{\mathsf{W}^{-1}}$. Averages over $\delta(\xi - \xi^*)\, e^{-\beta^* V}$ without this factor would require an additional reweighting by $|\nabla\xi|_{\mathsf{W}^{-1}}^{-1}$, which is not applied.

Reported quantities. The run reports the normalized weights, the Kish fraction $\mathrm{ESS}/n = (n\sum_i w_i^2)^{-1}$[44], the weighted mean midpoint energy and the weighted root-mean-square deviation of $\xi(\mathbf{M})$ from ξ^*. The last quantity measures the width of the window and so indicates how far the run is from the stiff limit.

6. COMMITTOR-LEARNING FORMULATION

This section gives the physical and statistical basis of the committor procedure in Methods Section 4.5. The biased walks, shooting trajectories and validation chains all use the kernel of Methods Section 4.1 at τ^*. Their core sets are constructed from the fingerprints and progress coordinates defined in Methods Section 4.3.

6.1. Committor of the Monte Carlo chain

The core sets and finite-step hitting probability are defined in Methods Section 4.5. Supplementary Sections 3.4 and 3.5 derive the diffusion approximation and explain the inverse masses in the variational functional. These results do not remove finite-step error. Shooting follows the implemented chain and supplies its validation reference (Supplementary Section 8.4).

6.2. Kolmogorov bias

With $q = \sigma(sz)$, $\nabla q = s\,q(1-q)\nabla z$, so $\ln|\nabla_u q|^2 = 2\ln s + \ln|\nabla_u z|^2 + 2\ln q + 2\ln(1-q)$. Equation (12) regularizes each logarithm separately and shifts the first by $-\ln\epsilon$ so that it is non-negative[27]. The gradient is that of z, since the factor $q(1-q)$ is already present in the other two terms. Because the first bracketed term is non-negative and $\ln(q+\delta)+\ln(1-q+\delta) \geq \ln\delta$ to first order in δ, $V_{\mathrm{K}} \leq 2\lambda|\ln\delta|/\beta^*$.

Biased ensemble. Exponentiating equation (12) gives

$$e^{-\beta^* V_{\mathrm{K}}} = \left[\frac{|\nabla_u z_\theta|^2+\epsilon}{\epsilon}\right]^{\lambda} (q_\theta+\delta)^{2\lambda}(1-q_\theta+\delta)^{2\lambda}. \quad \text{(S.59)}$$

For $\epsilon, \delta \to 0$ and with $|\nabla_u q_\theta|^2 = s^2 q_\theta^2(1-q_\theta)^2|\nabla_u z_\theta|^2$, the right-hand side tends to $(s^2\epsilon)^{-\lambda}|\nabla_u q_\theta|^{2\lambda}$. For $\lambda = 1$ the walks without V_z thus sample

$$p_{\mathrm{K}}(\mathbf{R}) = \frac{|\nabla_u q_\theta(\mathbf{R})|^2 e^{-\beta^* V(\mathbf{R})}}{\displaystyle\int |\nabla_u q_\theta(\mathbf{R}')|^2 e^{-\beta^* V(\mathbf{R}')}\,\mathrm{d}\mathbf{R}'}, \quad \text{(S.60)}$$

which is the normalized integrand of $K[q_\theta]$ in equation (10). This observation underlies the Kolmogorov bias of Kang et al.[26]. Sampling concentrates where the current model places the transition region, which is where the functional is least constrained by unbiased data. By Supplementary Section 2.3, V_{K} and p_{K} are invariant under rigid motions, so the biased walks remain covered by Supplementary Section 2.4.

6.3. Adaptive bias along the latent

The grid of V_z spans the range of z_θ over all stored configurations, widened by 2 on each side, with 201 points. Its range is recomputed in every iteration because the loss fixes only the product sz, so that refits may stretch z. A deposition at z_0 adds $h\exp[-(z-z_0)^2/(2\times 0.25^2)]$, with $h = \frac{1}{2}k_{\mathrm{B}}T^*\exp[-V_z(z_0)/(k_{\mathrm{B}}T^*(\gamma_{\mathrm{wt}}-1))]$ and bias factor $\gamma_{\mathrm{wt}} = 10$[49].

6.4. Reweighting biased blocks

Block k holds N_k configurations sampled with frozen bias b_k, and the basin block has $b_0 = 0$. With reduced biases $u_k(\mathbf{R}) = \beta^* b_k(\mathbf{R})$, the MBAR free energies[36] solve

$$e^{-f_k} = \sum_n \frac{e^{-u_k(\mathbf{R}_n)}}{\sum_l N_l e^{f_l - u_l(\mathbf{R}_n)}}, \quad \text{(S.61)}$$

where f_k is the reduced free energy of block k, the sum over n runs over all stored configurations and $f_0 = 0$ fixes the additive constant. The equations are iterated until the largest change is below 10^{-10}, for at most 500 iterations. Under equilibrium and overlap assumptions, the reweighting estimate for sample n is

$$w_n = \frac{\left[\sum_l N_l e^{f_l - u_l(\mathbf{R}_n)}\right]^{-1}}{\sum_m \left[\sum_l N_l e^{f_l - u_l(\mathbf{R}_m)}\right]^{-1}}. \quad \text{(S.62)}$$

Weights above the percentile specified in Supplementary Table S.4 are capped before normalization. This regularization trades bias for variance reduction.

7. HYPERPARAMETER CONFIGURATION

The proposal widths, free-particle pair coupling and molecular meeting criterion follow MCTSS[6]. The neural architecture, sigmoid slope and Kolmogorov-bias regularizers follow Trizio et al.[27]; shooting-point selection and likelihood follow AIMMD[20]. The remaining budgets, tolerances and resolution settings are numerical heuristics. Supplementary Tables S.2–S.4 collect the core parameters. Detailed move, restraint and ensemble protocols are given in Supplementary Section 5; neural and bias protocols are specified in Supplementary Sections 6 and 8.

Defaults are shared across systems, with molecular/two-dimensional differences shown by slashes in the tables. Degree-of-freedom scaling sets quench and shooting budgets, τ_{x} sets proposal widths and ℓ, and T_{phys} sets τ^*. The inexpensive two-dimensional potentials use longer sampling walks, more labelled paths and chains, and the distinct meeting criterion in Supplementary Section 5.2.

The shooting-loss multiplier is $\Gamma = 1$ for molecules and 0 for the two-dimensional models. Thus, two-dimensional shots contribute to diagnostics but not to refitting. No setting was tuned separately for an individual system; optional quench-budget and patience multipliers remained at their defaults.

TABLE S.2. **Exploration and continuation parameters.** Slash-separated values denote molecules / two-dimensional models; single values are shared. $\ell = \sqrt{3\hbar\tau_x/2}$ sets geometric resolution and L is the link length. The rung join fraction selects an exploration scale, not a benchmark success probability. Local-ensemble settings are given in Supplementary Section 5.9, and seed conventions in Supplementary Section 9.9.

Parameter	Value
Initial τ_0; rung ratio	1 fs; 2
Runs per rung; maximum rungs; halvings	4; 8; 3
Rung join fraction; steps per run	0.5; 5000
Pair coupling τ_{12}/τ_b	3 / 1
Meeting threshold for $d^2/(\hbar\tau_b)$	$3N/2$ / 0.5
Labelled paths; chains per state pair	8 / 12; 4 / 8
Slices; restraint κ_0	24; 1 eV
Slice tolerance; restraint stiffening	$\Delta\xi/4$; $\times 4$, up to 3 times
Maximum node gap; channel radius	ℓ; 3ℓ
Interior link points	$\max(4, \lceil 4L/\ell \rceil)$
Crossing votes; link bisections	3; 4

TABLE S.3. **Quench annealing schedules.** Steps are attempted moves per temperature level before multiplication by $r = \max(1, d/18)$. Endpoint, node and ramp thermal energies are also divided by r; label and chain-end levels are unscaled. Patience is four blocks of 25 steps, except zero for the ramp, with a 0.1 meV improvement threshold. Proposal-width adaptation is specified in Supplementary Section 5.4.

Use	k_BT levels (meV)	Steps	Initial parameter
Endpoint relaxation	20, 8, 3, 1, 0.5	400	τ_0
Path and crossing labels	5, 2, 1, 0.5	150	$\tau_x/16$
Chain node	20, 8, 3, 1	250	τ_x initially, then carried over
Restraint ramp step	3	125	carried over
Chain ends	5, 2, 1, 0.5	300	$\tau_x/16$

TABLE S.4. **Committor-training and sampling parameters.** Slash-separated values denote molecules / two-dimensional models. Biases are frozen before production; reweighting and weight capping precede refitting. Shooting diagnostics are evaluated separately from refitting. Network architectures and shot limits are given in Supplementary Table S.6, and deposition and storage intervals in Supplementary Section 8.3.

Parameter	Value
Iterations after initial fit	6
Initial basin steps per minimum	4000 / 20000
Biased exploration; production steps per walker	1500 / 5000; 12000 / 20000
Kolmogorov bias λ; ϵ, δ	1; 10^{-7}
Latent-bias width; initial height; bias factor	0.25; $k_BT^*/2$; 10
Epochs; minibatch; learning rate	300; 256 / 512; 10^{-3}
Loss weights α; Γ	1; 1 / 0
Reweighting cap	99th percentile
Shooting points per iteration; shots per point	16; 2
Probe shots per crossing	20

8. NEURAL COMMITTOR

This section specifies the model, loss and data-generation loop of Methods Section 4.5. Supplementary Algorithm 3 shows their execution order, Supplementary Table S.4 lists the shared settings, and Supplementary Table S.6 gives the system-dependent descriptors and shot budgets.

8.1. Network and descriptors

The descriptors are all $N(N-1)/2$ interatomic distances, each computed as $\sqrt{r_{ij}^2 + 10^{-12}}$, or the two coordinates of the two-dimensional models. The network maps them through two hidden layers of 32 units with hyperbolic-tangent activations to the scalar latent z_θ, with default PyTorch initialization, and $q_\theta = \sigma(3z_\theta)$. The same network instance is refitted in every iteration.

8.2. Loss and optimization

The mass-weighted gradient is obtained by automatic differentiation with respect to Cartesian coordinates, $|\nabla_u q|^2 = \sum_j m_j^{-1}(\partial q/\partial x_j)^2$. The first term of equation (11) normalizes weights within each uniformly shuffled minibatch. These updates need not be unbiased gradients of the full-data objective. The boundary and shooting terms use all core configurations and all committed shots, respectively. The shot term is

$$\ell_{\mathrm{shot}} = -\frac{\sum_k \left[n_B^{(k)} \ln q_\theta(\mathbf{R}_k) + n_A^{(k)} \ln\left(1 - q_\theta(\mathbf{R}_k)\right) \right]}{\sum_k \left(n_A^{(k)} + n_B^{(k)} \right)}, \tag{S.63}$$

evaluated through $\ln\sigma(\pm 3z)$ for numerical stability[20]. At the start of each fit, $\gamma = \Gamma\, \mathcal{L}_v^{(0)}/\ell_{\mathrm{shot}}^{(0)}$, where $\mathcal{L}_v^{(0)}$ is the

first term over the full data set. Without this scaling a few dozen shots would dominate the Boltzmann average of $|\nabla_u q|^2$, which is orders of magnitude smaller. Optimization uses Adam[59], with epochs, learning rate and minibatch sizes specified in Supplementary Table S.4. Data are shuffled from a fixed seed.

8.3. Iterative data generation

Supplementary Algorithm 3 gives the loop. The walkers are the two minima and the highest node of the lowest chain of every channel. In the exploration phase they take turns growing one shared V_z, depositing every 20 steps (50 in two dimensions). The bias is then frozen and every walker runs a production walk in parallel, storing every fifth configuration (every tenth in two dimensions). Walkers continue from their last configuration in the next iteration.

8.4. Shooting diagnostics

Candidates are the chain nodes and the configurations of the newest block. Their latents are histogrammed with $\min[30, \max(2, \lfloor n/10 \rfloor)]$ bins, and candidate n is drawn with the probability of equation (13), with the density $\rho(z_n)$ given by the count $h(z_n)$ of its bin. Sixteen points are drawn without replacement. From each, two independent unbiased chains at τ^* stop at the first accepted configuration inside a core. Their budgets are $2000\max(1, 3N/18)$ steps for molecules and 2000 steps in two dimensions. The implementation checks core membership only after an accepted move, so core-start shots need not satisfy the target boundary values. Before refitting, the current model is scored on these outcomes by ℓ_{shot} and by the ratio between the observed fraction of points whose two chains reached different states and its expectation $\langle 2q_\theta(1-q_\theta)\rangle$[20]. Both scores use only chains that committed.

Why these scores. For an initial point outside both cores and correctly implemented uncensored first-hitting outcomes, the two chains are independent and each reaches C_b first with probability $q(\mathbf{R})$ of equation (9). They reach different cores with probability

$$\Pr[\text{split} \mid \mathbf{R}] = q(1-q) + (1-q)q = 2q(\mathbf{R})\big[1 - q(\mathbf{R})\big], \tag{S.64}$$

so the ratio is one in expectation when $q_\theta = q$ on the shooting points and outcomes are uncensored. Conditioning on commitment within a finite budget can change this relation. For one point, the expected contribution of its chains to ℓ_{shot} is the cross-entropy $-q\ln q_\theta - (1-q)\ln(1-q_\theta)$. For uncensored outcomes, Gibbs' inequality makes this expectation minimal only at $q_\theta = q$, so the binomial likelihood is a strictly proper scoring rule[66]. A model that predicts $q_\theta = 1/2$ everywhere scores $\ell_{\text{shot}} = \ln 2$, which serves as the chance level.

Crossing probes report the complete-case fraction and binomial standard error defined in Methods Section 4.5, using the count in Supplementary Table S.4. Unresolved shots are counted separately and excluded from p_B. This fraction need not estimate the unconditional hitting probability if censoring depends on the eventual outcome.

9. SIMULATION DETAILS

This section records the energy function, endpoints, committor settings and independent evaluation references for each benchmark in Results Section 2.1. Supplementary Table S.5 summarizes the systems, and Supplementary Table S.6 lists the settings that vary with system size. All remaining settings use the defaults in Supplementary Section 7.

The three two-dimensional models share the following details. A single particle moves in the (x, y) plane, with coordinates in Å. Both coordinates are displaced at every step, and no frame is removed. The endpoints are fixed guesses of the two minima, relaxed by the endpoint quench of Supplementary Table S.3. Exploration runs also evaluate trials rejected by the pre-screen, for diagnostics only, and these evaluations are not counted as cost. The committor network reads (x, y) and has the architecture 2–32–32–1. Reference minima and first-order saddles are polished from fixed guesses with analytic gradients and checked with the Hessian. The reference committor is the finite-volume solution described in Supplementary Section 9.7, on the grid given for each model.

The molecular systems share the following details. One calculator instance is kept per walker, so that each self-consistent field calculation restarts from the walker's own previous solution. GFN2-xTB[34] is evaluated in process with tblite at accuracy 1, electronic temperature 300 K and at most 250 self-consistent iterations. The committor network reads all $N(N-1)/2$ interatomic distances. The endpoints are relaxed by the endpoint quench before exploration, so that non-optimized input structures are accepted.

9.1. Proton transfer coupled to a harmonic bath

Energy function. The model of MCTSS[6] is a symmetric double well along x, coupled bilinearly to a harmonic bath coordinate y,

$$\begin{aligned} V(x, y) = {} & \frac{V_{\text{BH}}}{x_0^4}(x - x_0)^2(x + x_0)^2 \\ & + \tfrac{1}{2} f y^2 - \alpha f x y, \end{aligned} \tag{S.65}$$

where $V_{\text{BH}} = 2000$ cm^{-1} is the uncoupled double-well barrier parameter and $x_0 = 0.5$ Å its uncoupled minimum coordinate, $f = 4\pi^2 m c^2 \tilde{\nu}^2$ the force constant of the bath for the wavenumber $\tilde{\nu} = 1000$ cm^{-1}, $\alpha = 0.7$ the coupling strength and $m = 1$ amu the particle mass. Energies are in cm^{-1} and no walls are added.

Simulation details. The endpoints are $(\mp x_0, 0)$. The physical parameter is $\tau^* = 10$ fs, which places the barrier near $8\,k_{\mathrm{B}}T^*$. Two minima share a state when they lie within 0.22 Å of each other.

Reference quantities. The bilinear coupling shifts the minima to approximately $(x, y) = \pm(0.603, 0.422)$ Å. The origin is the only saddle, with an actual barrier of about 4229 cm^{-1} relative to those coupled minima. The reference committor is computed on a 161×121 grid over $-1.2 \le x \le 1.2$ Å and $-0.8 \le y \le 0.8$ Å.

9.2. Müller–Brown potential

Energy function. The Müller–Brown potential[31] is used in the scaled form of Trizio et al.[27],

$$\begin{aligned} V(x,y) &= 0.15\Big[146.7 + \sum_{k=1}^{4} A_k e^{g_k(x,y)}\Big], \\ g_k(x,y) &= a_k(x-x_k)^2 + b_k(x-x_k)(y-y_k) \\ &\quad + c_k(y-y_k)^2, \end{aligned} \tag{S.66}$$

where, for $k = 1$ to 4, A_k is -200, -100, -170 and 15; a_k is -1, -1, -6.5 and 0.7; b_k is 0, 0, 11 and 0.6; c_k is -10, -10, -6.5 and 0.7; x_k is 1, 0, -0.5 and -1; and y_k is 0, 0.5, 1.5 and 1. Energies are in units of $k_{\mathrm{B}}T$ at 300 K. The exponent g_k is capped at 50 far outside the walls. Harmonic walls $1000[\min(x+1.3, 0)^2 + \max(x-1.0, 0)^2]$ confine x, as in the reference simulations[27].

Simulation details. The particle mass is 4 amu and $T_{\mathrm{phys}} = 300$ K, so $\tau^* = 25.46$ fs. The endpoints are the guesses $(-0.558, 1.442)$ for A and $(0.623, 0.028)$ for B. The surface has a third minimum near $(-0.050, 0.467)$, which the search must find as an intermediate. Two minima share a state when they lie within 0.15 Å of each other.

Reference quantities. The two saddles are polished from guesses near $(-0.822, 0.624)$ and $(0.212, 0.293)$. The reference committor is computed on a 181×181 grid over $-1.5 \le x \le 1.2$ and $-0.5 \le y \le 2.2$.

9.3. Two-channel potential

Energy function. The two-path potential of Trizio et al.[27] is a quartic double well with two channels between the basins, the upper one raised and the lower one deepened by Gaussians,

$$\begin{aligned} V(x,y) &= 10\big[2 + \tfrac{4}{3}x^4 - 2y^2 + y^4\big] + \tfrac{100}{3}x^2(y^2-1) \\ &\quad + \sum_{k=1}^{3} A_k \exp\Big[-\frac{(x-x_k)^2 + (y-y_k)^2}{0.4^2}\Big], \end{aligned} \tag{S.67}$$

where the Gaussian amplitudes are $(A_1, A_2, A_3) = (7, 1, -6)$ and their centres (x_k, y_k) are $(-0.7, 0.8)$, $(1.0, -0.3)$ and $(-1.0, -0.6)$. Energies are in units of $k_{\mathrm{B}}T$ at 300 K, and harmonic walls $1000[\min(x+2, 0)^2 + \max(x-2, 0)^2]$ confine x.

Simulation details. The particle mass, T_{phys} and τ^* are those of the Müller–Brown potential. The endpoints are the guesses $(-1.033, -0.350)$ for A and $(1.122, 0.043)$ for B. Two minima share a state when they lie within 0.20 Å of each other. The test of this system is whether both channels are found and certified as separate channels (Methods Section 4.3).

Reference quantities. The saddle of the upper channel is polished from a guess near $(-0.424, 0.637)$ and that of the lower channel, which is almost flat, from a guess near $(0.075, -0.995)$. The reference committor is computed on a 181×181 grid over $-2 \le x, y \le 2$.

9.4. $\mathrm{S_N2}$ reaction $\mathrm{Cl^-} + \mathrm{CH_3Br} \rightarrow \mathrm{CH_3Cl} + \mathrm{Br^-}$

Energy function. The benchmark in Results uses GFN2-xTB energies and the corresponding independently refined reference saddle. The system has six atoms, charge -1 and multiplicity 1.

Simulation details. The supplied endpoints were optimized at the HF-3c level, then relaxed by the energy-only endpoint quench on GFN2-xTB. States are bond graphs, and ξ_{bond} reduces to $r(\mathrm{C{-}Br}) - r(\mathrm{C{-}Cl})$. $T_{\mathrm{phys}} = 298.15$ K. The input-file calculator default must be overridden when reproducing the GFN2-xTB comparison.

Reference quantities. The molecular reference saddle was refined on GFN2-xTB for comparison only, as described in Results. Energy differences in Table I use that same-surface saddle. The published HF-3c barriers of 5.5 and 143.5 $\mathrm{kJ\,mol^{-1}}$[6] refer to a different energy surface and are not used to assess these GFN2-xTB results.

9.5. Gas-phase reactions of the Baker test set

Energy function. Eight reactions from the standard transition-state test set of Baker and Chan[32], in the geometries used by Suresh and Vogiatzis[33], are computed with GFN2-xTB. Their charges and multiplicities are listed in Supplementary Table S.5.

Reactions. Supplementary Table S.5 lists all eight transformations with their sizes, charges and multiplicities. They cover closed-shell and radical hydrogen shifts, keto–enol tautomerization, dissociation, addition and radical ring opening. These classes test distinct changes in the bond-graph fingerprint rather than eight repetitions of one reaction type.

Simulation details. The endpoints are those of Suresh and Vogiatzis[33] and were not optimized. States are bond graphs, and $T_{\mathrm{phys}} = 298.15$ K for all eight reactions. Each reaction therefore tests the search from unrelaxed input structures on a fast semiempirical oracle.

Reference quantities. The reference structures are the transition states reported by Suresh and Vogiatzis,

located with xtb and the Berny optimizer[33]. For this benchmark, these geometries initialized independent Sella reference refinement on GFN2-xTB, followed by a finite-difference Hessian check. Only these refined same-surface references are used for the saddle comparisons in Results; the input guesses are evaluated separately.

9.6. Alanine dipeptide

Energy function. Alanine dipeptide (Ace-Ala-Nme, 22 atoms) in vacuum is described by an energy-only reimplementation of AMBER ff99SB[35]. Its parameters were exported from the GROMACS run input of Trizio et al.[27]. The energy contains harmonic bonds, including the X–H bonds, harmonic angles, periodic proper and improper dihedrals, scaled 1–4 Lennard-Jones and Coulomb terms, and all remaining non-bonded pairs without cutoff. The calculator returns energies only.

Simulation details. The endpoints are the starting frames of the $C7_{eq}$ and $C7_{ax}$ simulations of the same reference and were not optimized. States are defined by (ϕ, ψ), with ϕ through atoms 5–7–9–15 and ψ through atoms 7–9–15–17 in the GROMACS order. Two configurations share a state when both periodic differences are at most 60°. New states are disallowed, so every quench is assigned to $C7_{eq}$ or $C7_{ax}$. $T_{phys} = 300$ K, as in the reference simulations.

Committor model details. The network reads all 231 interatomic distances. This differs from the 45 heavy-atom distances used by Trizio et al.[27], because the descriptor layer of tReax_expedition is the same for every molecule.

Reference quantities. The reference is a first-order saddle of the same force field near $(\phi, \psi) = (-2.1°, -26.4°)$, located independently of the search, with a barrier that depends on the reactant energy convention. Relative to the lowest state-A minimum recorded by the retained search, the barrier is 33.46 kJ mol^{-1}; the earlier 36.4 kJ mol^{-1} value used a different minimum. Results primarily compares absolute saddle-referenced differences to avoid this offset. The reference serves only for evaluation.

9.7. Evaluation against reference quantities

None of the following enters the search. For the two-dimensional models, the reference committor is the finite-volume solution of $\nabla \cdot (e^{-\beta^* V} \nabla q) = 0$ on the grids of Supplementary Sections 9.1 to 9.3, with harmonic means of $e^{-\beta^* V}$ on cell faces. The outer boundaries reflect, and discs around the exact minima nearest to the states found by the search absorb. The discs have the basin radii of Supplementary Table S.5, and $\beta^*(V - V_{min})$ is capped at 300 to avoid underflow. The learned committor is compared on grid points with $0.05 < q < 0.95$ and $V \leq V_S + 3k_B T^*$, where V_S is the lowest saddle, and by a Boltzmann-weighted error. The grid solution is a numerical reference for overdamped diffusion, not an exact finite-step Monte Carlo committor. Its absorbing discs also differ from the core rules used by shooting, so their comparison includes both dynamical and boundary-definition differences. For molecules, the Results comparisons use independently refined saddle structures on the same oracle, checked for exactly one imaginary mode, and compare their energies with E_{up}. Potential-energy gradients and finite-difference Hessians are used only in this reference calculation, not in the search. Published barriers are compared only when the level of theory matches.

9.8. Barrier estimates from single coupled-walker runs

For comparison with the estimator of MCTSS[6], a single exploration run can be analysed without labelling, chain growth or committor learning. Either the pooled samples of both walkers or the weighted midpoints of equation (6) are binned into 15 equal-width bins of the progress p. The forward and backward barriers are the highest bin mean minus the mean of the reactant and product end bins, with a within-run spread from 200 resamples of individual stored frames. This framewise resampling ignores serial dependence and is not calibrated independent-search uncertainty. Walker samples belong to no single temperature, because each walker samples τ_b while apart, the midpoint $2\tau_b$ when together, and the approach is not stationary. The midpoint weighting identity applies to stationary sampling of Π at $\tau_{mid} = 2\tau_b$ (Supplementary Section 2.5). An exploration run stopped at a join does not establish stationarity. Bin means lie above the zero-temperature profile by roughly $\frac{1}{2}(d-1)k_B T$, and a weighted linear fit in $1/\tau$ over at least three runs extrapolates the estimate to zero effective temperature. These estimators are heuristics and are reported separately from E_{up}.

9.9. Software, random numbers and reproducibility

The implementation is written in Python 3.12 with NumPy, SciPy, ASE[56], PyTorch[58] and tblite. The reported hardware is specified in Methods Section 4.6. Exact package versions for the reported calculations are unavailable. Monte Carlo moves draw from the Mersenne Twister generator of the Python standard library, seeded per task from the base seed s ($s = 1$ by default). Run k of the exploration ladder uses seed $1000s + k$, counted across rungs. Labelled path k uses $2000s + k$, chain k of a state pair uses $3000s + k$, and transition-state ensemble job j uses $3000s + j$. Committor learning seeds PyTorch with s and draws the seeds of its walks and shots from a NumPy generator seeded with s. Tasks run in a pool of spawned worker processes, four by default, each with one thread and double-precision tensors. Every task carries its own seed and results are gathered in submission order. This design reduces scheduling dependence, but

does not establish reproducibility across software revisions or numerical libraries. Energies from self-consistent field calculations depend on the restart guess within the convergence threshold, so runs on ORCA or tblite are reproducible up to that tolerance. The implementation can write settings, state records, chain energies, costs and structures, with checkpoints and ensemble samples for optional stages.

10. TREAX_EXPEDITION ALGORITHM

Supplementary Algorithm 1 gives the complete search between two endpoints and places the optional committor and transition-state ensemble steps in order. Supplementary Algorithm 2 details the growth and certification of one continuation chain, and Supplementary Algorithm 3 the iterative learning of the committor for one pair of states. Line-level settings refer to Supplementary Tables S.2–S.4.

Supplementary Algorithm 1. tReax_expedition search between two endpoints.

Require: endpoints $\mathbf{R}_A^0, \mathbf{R}_B^0$; oracle V; fingerprint; settings
relax both endpoints by quenching; register states A and B; stop if they coincide
$\tau \leftarrow \tau_0$; run four coupled-walker runs at τ; halve τ (at most three times) until at least half join
repeat $\tau_\mathrm{x} \leftarrow \tau$; $\tau \leftarrow 2\tau$; run four coupled-walker runs at τ
until fewer than half join, or eight rungs
pick up to eight joined paths with $\tau_\mathrm{b} \le \tau_\mathrm{x}$; label each by bisection with quenches
for all state pairs (S_a, S_b) connected by a label change **do**
  build ξ; choose up to four seeds by farthest-point selection
  grow a chain from every seed (Algorithm 2); keep certified chains
  if no chain is certified and committor learning is requested **then**
    learn q_θ from the seeds (Algorithm 3); regrow the chains along z_θ
  end if
  group chains into channels; keep the lowest chain of each
end for
on each channel, bisect the three-vote label along the chain; refine the link four times
optional: learn q_θ at τ^* for each pair (Algorithm 3); shoot from the crossings
optional: sample the straddling pair at $\tau_\mathrm{b} = \tau^*/2$ on each channel
report the minimax route from A to B, all E_up, crossings and oracle calls per step

Supplementary Algorithm 2. One restrained continuation chain between S_a and S_b.

Require: seed $\mathbf{R}_0$; coordinate ξ; minima of S_a, S_b; τ_x
slice grid $\xi_0, \ldots, \xi_{24}$ between the minima; $\kappa \leftarrow \kappa_0/\Delta\xi^2$; $\ell \leftarrow \sqrt{3\hbar\tau_\mathrm{x}/2}$
seed node: restrained quench at clip$(\xi(\mathbf{R}_0))$
for all sweep directions, slices beyond $\Delta\xi/4$ from the seed, in order **do**
  move the restraint centre to ξ_k in five short quenches from the previous node
  quench at ξ_k; quadruple κ and repeat while the miss exceeds $\Delta\xi/4$ (at most three times)
  record the highest energy held along these quenches as E_k^traj
  insert nodes at half spacing while two nodes are farther apart than ℓ (depth at most three)
end for
quench both end nodes without restraint; register the states they reach; record E^desc
evaluate V on every straight link at $\max(4, \lceil 4L/\ell \rceil)$ interior points
$E_\mathrm{up} \leftarrow$ equation (8); certified $\leftarrow$ ends reach (S_a, S_b) and $E_\mathrm{up} < \infty$

Supplementary Algorithm 3. Committor learning for one state pair.

Require: minima of S_a, S_b; chain tops and nodes; quench-label crossings; τ^*
unbiased walks from both minima; core configurations define the boundary sets; store the block with $b_0 = 0$
fit q_θ on the stored data (MBAR weights)
for $k = 1, \ldots, 6$ **do**
  $V_\mathrm{bias} \leftarrow V_\mathrm{K}[q_\theta] + V_z$, with the V_z grid sized from the stored latents
  walkers take turns depositing V_z; freeze V_bias
  production walks under the frozen bias; store the block with its bias
  draw 16 shooting points (Lorentzian in z_θ over density); run two unbiased chains from each
  score the current model on the new shots; add the shots to the data
  refit q_θ with MBAR weights capped at the 99th percentile
end for
shoot 20 chains from each quench-label crossing; report p_B against q_θ

TABLE S.5. **Energy surfaces and state definitions used in the benchmark.** Each row specifies the energy oracle, system size, charge and spin multiplicity, physical-temperature parameter and fingerprint used for quench-based state assignment. Radii denote the two-dimensional basin-matching tolerance; the alanine tolerance applies separately to both wrapped backbone dihedrals. All gas-phase results use GFN2-xTB, including S_N2. The search used the same molecular defaults, with the stated two-dimensional exceptions.

System	Energy	Size	Charge, mult.	T_{phys} (τ^*)	State fingerprint
Proton transfer with bath[6]	analytic (cm^{-1})	1 particle, 1 amu	–	$\tau^* = 10$ fs	region, 0.22 Å
Müller–Brown, scaled[27, 31]	analytic (reduced)	1 particle, 4 amu	–	300 K (25.46 fs)	region, 0.15 Å
Two-channel[27]	analytic (reduced)	1 particle, 4 amu	–	300 K (25.46 fs)	region, 0.20 Å
$Cl^- + CH_3Br$ (S_N2)	GFN2-xTB	6 atoms	−1, 1	298.15 K	bond graph
$HCN \rightarrow HNC$	GFN2-xTB	3	0, 1	298.15 K	bond graph
$HCCH \rightarrow H_2CC$	GFN2-xTB	4	0, 1	298.15 K	bond graph
$HCOH \rightarrow H_2CO$	GFN2-xTB	4	0, 1	298.15 K	bond graph
$H_2CO \rightarrow H_2 + CO$	GFN2-xTB	4	0, 1	298.15 K	bond graph
$CH_3O^\bullet \rightarrow {}^\bullet CH_2OH$	GFN2-xTB	5	0, 2	298.15 K	bond graph
$CH_3CHO \rightarrow CH_2CHOH$	GFN2-xTB	7	0, 1	298.15 K	bond graph
$C_2H_4 + HF \rightarrow CH_3CH_2F$	GFN2-xTB	8	0, 1	298.15 K	bond graph
$C_3H_5O_2^\bullet$ ring opening	GFN2-xTB	10	0, 2	298.15 K	bond graph
Alanine dipeptide, $C7_{\mathrm{eq}} \rightarrow C7_{\mathrm{ax}}$	ff99SB, vacuum	22	0, 1	300 K	(ϕ, ψ), 60°

TABLE S.6. **System-dependent neural descriptors and shooting budgets.** Molecular networks receive all $N(N-1)/2$ interatomic distances; the two-dimensional models receive (x, y). Network sizes list the input, hidden and output dimensions. Γ controls the binomial shooting-loss contribution, and minibatch denotes configurations per optimization batch. A shot stops on reaching a core or exhausting $2000\max(1, 3N/18)$ steps, rounded for molecules. Non-committing shots are recorded separately and excluded from the reported product fraction.

System	Atoms	Descriptors	Network	Γ, minibatch	Shot budget
Two-dimensional models (3)	–	2	2–32–32–1	0, 512	2000
$HCN \rightarrow HNC$	3	3	3–32–32–1	1, 256	2000
HCCH, HCOH, H_2CO reactions	4	6	6–32–32–1	1, 256	2000
$CH_3O^\bullet \rightarrow {}^\bullet CH_2OH$	5	10	10–32–32–1	1, 256	2000
$Cl^- + CH_3Br$ (S_N2)	6	15	15–32–32–1	1, 256	2000
$CH_3CHO \rightarrow CH_2CHOH$	7	21	21–32–32–1	1, 256	2333
$C_2H_4 + HF \rightarrow CH_3CH_2F$	8	28	28–32–32–1	1, 256	2667
$C_3H_5O_2^\bullet$ ring opening	10	45	45–32–32–1	1, 256	3333
Alanine dipeptide	22	231	231–32–32–1	1, 256	7333

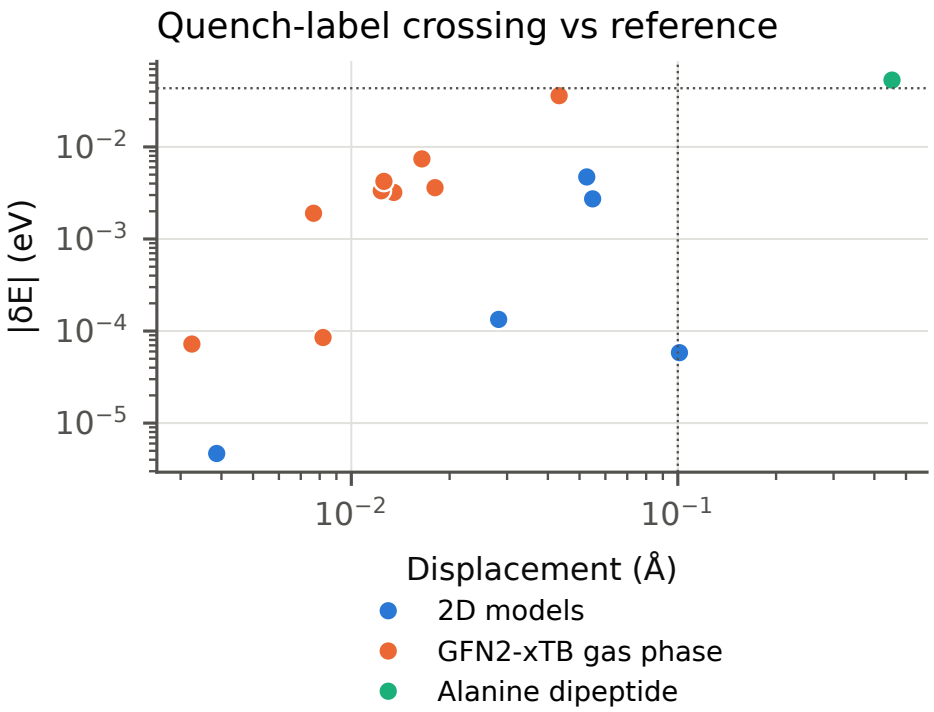


FIG. S.1. **Crossing displacement and energy deviation measure different aspects of localization.** Points compare quench-label crossings with the corresponding same-surface reference saddles. Blue denotes the five two-dimensional saddle comparisons, orange the nine primary gas-phase channels and green the lowest alanine dipeptide chain. The molecular displacement is symmetry-adjusted all-atom RMSD; two-dimensional displacements are Euclidean distances. Both axes are logarithmic; $|\delta E| = |V(\mathbf{R}_0^{\ddagger}) - E_{\mathrm{reference}}|$ is the absolute crossing-energy deviation. Dotted guides mark 0.1 Å and approximately 1 kcal mol^{-1}; they are visual scales, not statistical thresholds. The 15 points derive from one search per system, with no uncertainty bars or cross-system correlation fit.

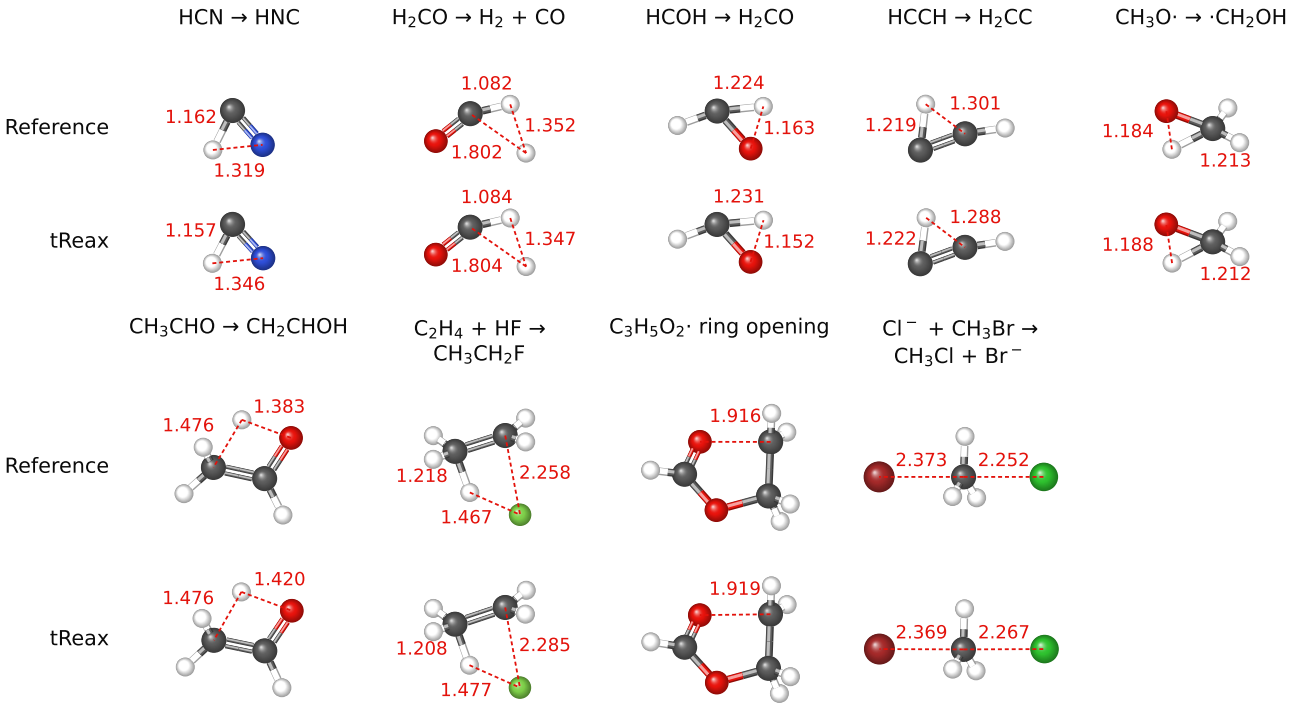


FIG. S.2. **Local geometric comparison of gas-phase reference saddles and quench-label crossings.** Each reaction is shown as a pair, with the independently refined GFN2-xTB first-order saddle (Reference) above the primary tReax_expedition crossing (tReax). The upper group shows HCN isomerization, formaldehyde dissociation, HCOH isomerization, acetylene rearrangement and methoxy isomerization, from left to right; the lower group shows acetaldehyde tautomerization, HF addition to ethylene, radical ring opening and S_N2 substitution. Red dashed segments identify selected interatomic distances, labelled in Å. Carbon, hydrogen, oxygen and nitrogen are grey, white, red and blue; fluorine, chlorine and bromine are light green, green and dark red. These local distances complement the all-atom RMSDs in Table I and the energy profiles in main Fig. 2; they do not establish stationarity or a finite-temperature isocommittor for the crossing structures. One retained crossing is shown per reaction.

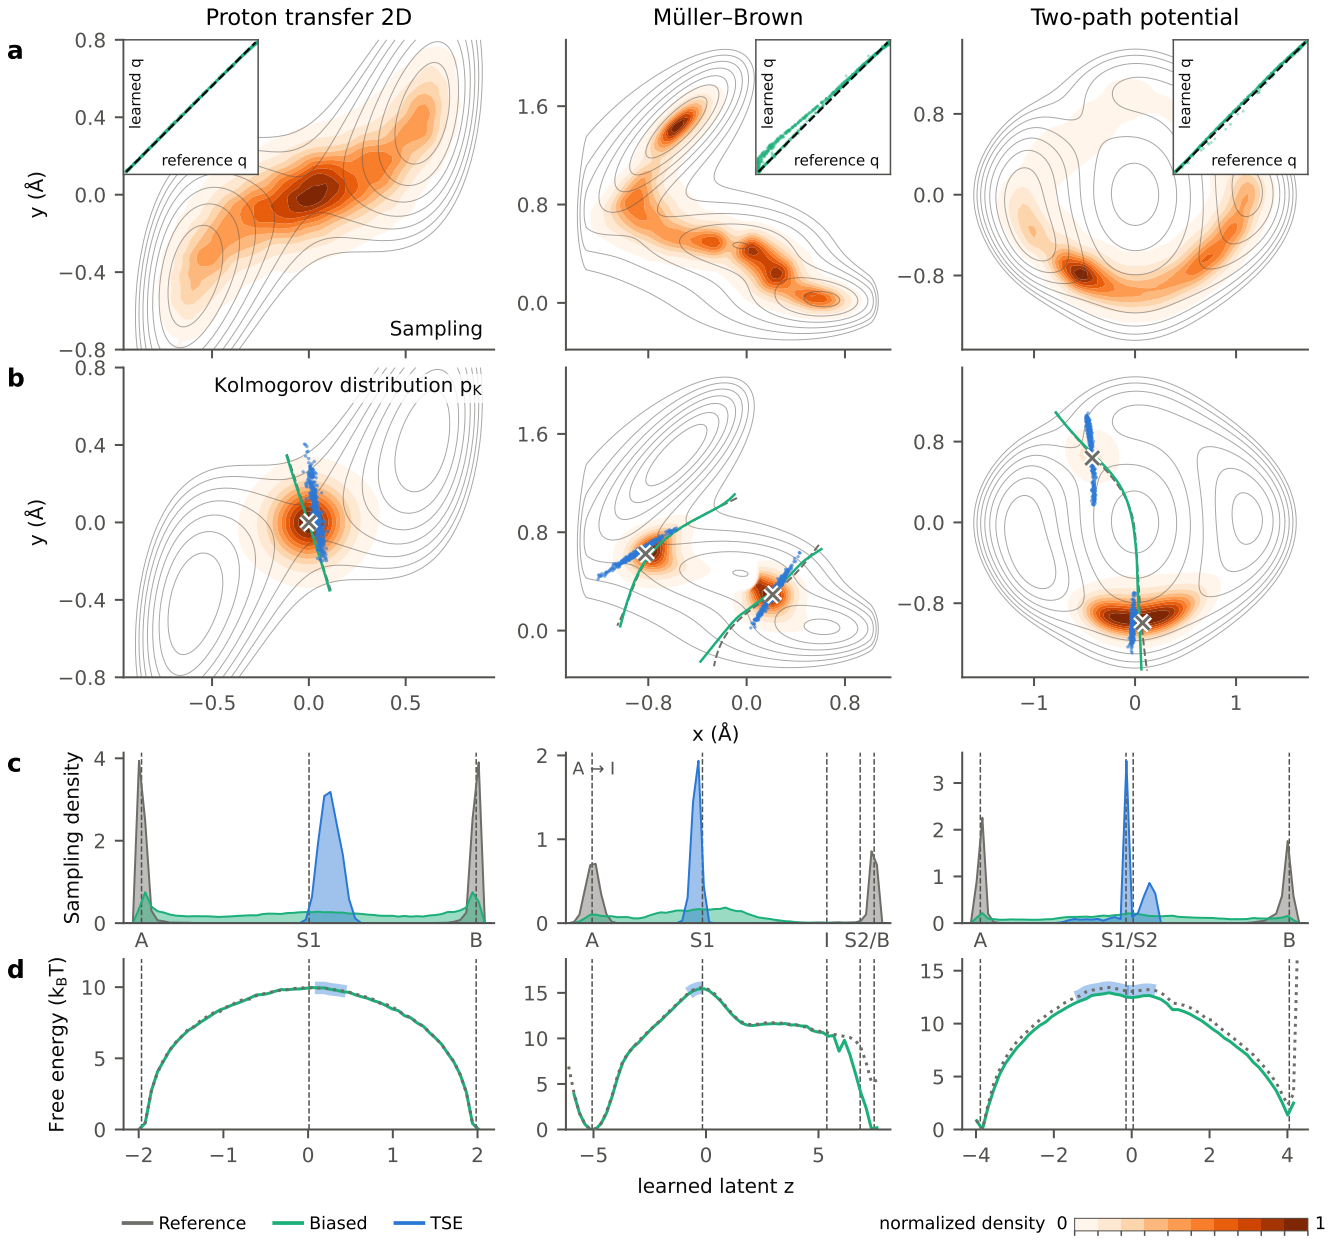


FIG. S.3. **Broad biased sampling and local transition-region sampling serve different purposes.** The panels illustrate sampling support and latent projections rather than quantitative committor or free-energy benchmarks. Columns show proton transfer, Müller–Brown and the two-path potential. **a**, Two-dimensional biased sampling densities; insets compare the learned committor with the numerical diffusion reference, with dashed diagonals marking equality. **b**, Kolmogorov density p_K (orange; equation (S.60)), local straddling-pair midpoint samples (TSE, blue), learned $q_\theta = 1/2$ contours (green), numerical reference $q = 1/2$ contours (grey dashed) and reference saddles (crosses). Grey background contours in a,b show potential energy. Orange maps use the normalized-density colour scale shown. **c**, Distributions along the learned latent z: numerical equilibrium reference (grey), biased training samples (green) and TSE midpoints (blue). Vertical guides locate the labelled minima and saddles. Both Müller–Brown projections in c,d refer to the indicated $A \to I$ committor; b includes both saddle neighbourhoods. **d**, Latent free-energy projections in k_BT units from the numerical reference (grey) and MBAR-reweighted biased samples (green), up to an additive constant. Blue segments mark the latent intervals occupied by TSE samples on the reference profile; they are neither TSE-derived free-energy curves nor confidence intervals. Biased walkers start from the two minima and retained channel tops, combining V_K with V_z before frozen-bias production. The TSE pair straddles a surface through the quench-label crossing without using q_θ in its side constraint. Its restricted support does not determine a full latent free-energy profile.