

1

Supporting Information: MAPLE: A

Machine-Learning Force-FieldNative Platform

2

for Automated Reaction Modeling and Enzyme

Design

3 Details for Algorithmic Developments

4 Batched L-BFGS Optimization with Dynamic Batch Shrinking

5 For large-scale geometry optimizations involving multiple independent molecular struc-
6 tures, we developed a GPU-accelerated batched L-BFGS optimizer that exploits parallelism
7 across structures while maintaining individual quasi-Newton histories for each system. This
8 approach significantly reduces wall-clock time for high-throughput computational screening
9 workflows.

10 The Limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) algorithm¹ is a quasi-
11 Newton method that approximates the inverse Hessian using a limited history of gradient and
12 position changes. In standard L-BFGS, the search direction \mathbf{p}_k at iteration k is computed via
13 the two-loop recursion algorithm, which requires storing the most recent m pairs of vectors
14 $\{\mathbf{s}_i, \mathbf{y}_i\}_{i=k-m}^{k-1}$, where $\mathbf{s}_i = \mathbf{x}_{i+1} - \mathbf{x}_i$ is the position change and $\mathbf{y}_i = \nabla f_{i+1} - \nabla f_i$ is the
15 gradient change. Our implementation extends this to optimize B structures simultaneously on
16 GPU with three key innovations.

17 **Fixed-dimension tensor batching.** All molecular structures are represented in a unified
18 coordinate tensor of shape (B, n_{\max}) , where $n_{\max} = 3N_{\max}$ is fixed at initialization based on
19 the largest structure in the batch. Smaller structures are zero-padded, and a binary mask tensor
20 $\mathbf{M} \in \{0, 1\}^{B \times n_{\max}}$ identifies valid atomic coordinates. This uniform dimensionality enables
21 efficient batched linear algebra operations (element-wise products, dot products, vector addi-
22 tions) to be executed in parallel across all structures with a single GPU kernel launch, avoiding
23 the overhead of sequential per-structure processing.

24 **Independent per-batch L-BFGS histories.** Each structure maintains its own quasi-Newton
25 memory consisting of m stored vector pairs and curvature information. The history for batch
26 b is represented as $\{\mathbf{S}^{(b)}, \mathbf{Y}^{(b)}, \boldsymbol{\rho}^{(b)}\}$, where $\mathbf{S}^{(b)}, \mathbf{Y}^{(b)} \in \mathbb{R}^{m \times n_{\max}}$ store the m most recent \mathbf{s} and
27 \mathbf{y} vectors, and $\boldsymbol{\rho}^{(b)} \in \mathbb{R}^m$ stores the corresponding inverse curvatures $\rho_i = 1/(\mathbf{y}_i^T \mathbf{s}_i)$. In our
28 batched implementation, these are stored as tensors $\mathbf{S} \in \mathbb{R}^{B \times m \times n_{\max}}$, $\mathbf{Y} \in \mathbb{R}^{B \times m \times n_{\max}}$, and

29 $\boldsymbol{\rho} \in \mathbb{R}^{B \times m}$. The two-loop recursion is executed in parallel for all batches:

$$\begin{aligned}
& \text{Backward pass: } \mathbf{q}^{(b)} \leftarrow \nabla f_k^{(b)}, \\
& \text{for } i = k - 1, \dots, k - m : \quad \alpha_i^{(b)} = \rho_i^{(b)} (\mathbf{s}_i^{(b)} \cdot \mathbf{q}^{(b)}), \quad \mathbf{q}^{(b)} \leftarrow \mathbf{q}^{(b)} - \alpha_i^{(b)} \mathbf{y}_i^{(b)} \\
& \text{Scaling: } \quad \gamma^{(b)} = (\mathbf{y}_{k-1}^{(b)} \cdot \mathbf{s}_{k-1}^{(b)}) / (\mathbf{y}_{k-1}^{(b)} \cdot \mathbf{y}_{k-1}^{(b)}), \quad \mathbf{z}^{(b)} = \gamma^{(b)} \mathbf{q}^{(b)} \\
& \text{Forward pass: for } i = k - m, \dots, k - 1 : \\
& \quad \beta_i^{(b)} = \rho_i^{(b)} (\mathbf{y}_i^{(b)} \cdot \mathbf{z}^{(b)}), \quad \mathbf{z}^{(b)} \leftarrow \mathbf{z}^{(b)} + \mathbf{s}_i^{(b)} (\alpha_i^{(b)} - \beta_i^{(b)}) \\
& \text{Search direction: } \quad \mathbf{p}_k^{(b)} = -\mathbf{z}^{(b)}
\end{aligned} \tag{1}$$

30 All dot products and vector updates are computed using masked batched operations, ensuring
31 that padding does not affect optimization trajectories. Each structure independently accumu-
32 lates curvature information appropriate to its own potential energy surface topology.

33 **Dynamic batch shrinking with history preservation.** Structures converge at different
34 rates depending on initial geometry quality and PES characteristics. Rather than waiting for all
35 structures to converge, our optimizer dynamically removes converged structures from the batch
36 at each iteration. When a subset of structures converges, the batch dimension B is reduced
37 via index selection, and all history tensors $\{\mathbf{S}, \mathbf{Y}, \boldsymbol{\rho}\}$ are correspondingly sliced to retain only
38 unconverged structures:

$$B \rightarrow B' = |\{\text{unconverged structures}\}|, \quad \mathbf{S} \rightarrow \mathbf{S}[\text{survive_idx}], \quad \text{etc.} \tag{2}$$

39 This shrinking operation maintains GPU occupancy by reducing wasted computation on struc-
40 tures that have already reached their optima.

41 The step size for each structure is independently clipped to prevent large displacements:

$$\mathbf{p}_k^{(b)} \leftarrow \mathbf{p}_k^{(b)} \cdot \min \left(1, \frac{\Delta_{\max}}{\max |\mathbf{p}_k^{(b)}|} \right) \tag{3}$$

42 where $\Delta_{\max} = 0.2 \text{ \AA}$ is the maximum allowed atomic displacement per step. The initial inverse
43 Hessian approximation is set to $H_0^{-1} = \gamma_0^{-1} \mathbf{I}$ with curvature parameter $\gamma_0 = 70 \text{ Eh/\AA}^2$, and
44 the history depth is $m = 5$ vector pairs. All optimizations used double-precision (float64)

45 arithmetic to ensure numerical stability in force and gradient evaluations.

46 **Dual-Shift Partitioned Rational Function Optimization**

47 In standard P-RFO,^{2,3} the optimization problem is partitioned into two subspaces: the tran-
 48 sition mode (uphill direction, eigenvalue $\lambda_- < 0$) where energy is maximized, and the remain-
 49 ing modes (downhill directions, $\lambda_+ > 0$) where energy is minimized. The step \mathbf{s} is constrained
 50 by a trust radius R via a single shift parameter α that scales both subspaces simultaneously to
 51 satisfy $\|\mathbf{s}\| \leq R$.

52 We developed a dual-shift variant (DS-P-RFO) that introduces independent control over
 53 each subspace. The total step \mathbf{s} is decomposed as:

$$\mathbf{s} = \mathbf{s}_- + \mathbf{s}_+ \quad (4)$$

54 where $\mathbf{s}_- \in \text{span}\{v_j\}$ represents the step in the uphill subspace (typically one-dimensional
 55 for first-order saddle points), and $\mathbf{s}_+ \in \text{span}\{v_i : i \neq j\}$ represents the step in the downhill
 56 subspace.

57 The key innovation is decomposing the global trust-region constraint:

$$\|\mathbf{s}\|^2 = \|\mathbf{s}_-\|^2 + \|\mathbf{s}_+\|^2 \leq R^2 \quad (5)$$

58 into subspace-specific constraints with adaptive allocation. We define:

$$R_-^2 = \alpha R^2, \quad R_+^2 = (1 - \alpha) R^2 \quad (6)$$

59 where $\alpha \in [0, 1]$ is determined by the relative magnitudes of unconstrained Newton steps in
 60 each subspace. Specifically, we compute:

$$\mathbf{s}_-^{\text{unc}} = -\frac{\mathbf{g}_-}{\lambda_-}, \quad \mathbf{s}_+^{\text{unc}} = -\frac{\mathbf{g}_+}{\lambda_+} \quad (7)$$

61 where \mathbf{g}_- and \mathbf{g}_+ are the gradient projections onto the uphill and downhill subspaces, respec-

62 tively. The allocation parameter is then:

$$\alpha = \frac{\|\mathbf{s}_-^{\text{unc}}\|^2}{\|\mathbf{s}_-^{\text{unc}}\|^2 + \|\mathbf{s}_+^{\text{unc}}\|^2} \quad (8)$$

63 with bounds $\alpha \in [0.05, 0.95]$ to prevent extreme allocations. This heuristic automatically dedi-
 64 cates larger trust-region budgets to subspaces with stronger gradient driving forces or shallower
 65 curvature.

66 For each subspace, we solve independently for shift parameters μ_- and μ_+ via bisection
 67 search. In the uphill subspace (maximization):

$$\sum_{i \in \mathcal{M}_-} \left(\frac{g_{-,i}}{\lambda_{-,i} - \mu_-} \right)^2 = R_-^2 \quad (9)$$

68 and in the downhill subspace (minimization):

$$\sum_{i \in \mathcal{M}_+} \left(\frac{g_{+,i}}{\lambda_{+,i} - \mu_+} \right)^2 = R_+^2 \quad (10)$$

69 where the sums are over eigenmode indices in each subspace. The step components are then:

$$s_{-,i} = -\frac{g_{-,i}}{\lambda_{-,i} - \mu_-}, \quad s_{+,i} = -\frac{g_{+,i}}{\lambda_{+,i} - \mu_+} \quad (11)$$

70 This approach eliminates the iterative micro-cycles required in traditional RS-P-RFO, offering
 71 improved computational efficiency while maintaining numerical stability through independent
 72 optimization of physically distinct objectives.

73 **Dimer Method with Hessian-Vector Products**

74 For saddle point searches without good initial guesses, we employed a modified dimer
 75 method⁴ that leverages automatic differentiation to avoid numerical instability from finite-
 76 difference approximations. The dimer method is a minimum-mode following algorithm that
 77 locates first-order saddle points using only gradient information, making it particularly suitable
 78 for large systems where full Hessian calculations are prohibitive.

79 The dimer consists of two images separated by distance 2Δ along a unit vector \mathbf{n} , positioned
80 at $\mathbf{R}_0 \pm \Delta\mathbf{n}$. The curvature along the dimer axis is approximated by:

$$\kappa = \mathbf{n}^T \mathbf{H} \mathbf{n} \approx \frac{\mathbf{F}^+ - \mathbf{F}^-}{2\Delta} \quad (12)$$

81 where \mathbf{F}^\pm are forces at the two dimer endpoints. Standard implementations use finite-difference
82 with $\Delta \approx 0.005 \text{ \AA}$, requiring two force evaluations per rotation step to estimate the Hessian-
83 vector product $\mathbf{H}\mathbf{n}$.

84 We developed an automatic differentiation (autograd) implementation that computes $\mathbf{H}\mathbf{n}$
85 analytically via forward-mode differentiation:

$$\mathbf{H}\mathbf{n} = \nabla_{\mathbf{R}} (\nabla_{\mathbf{R}} E(\mathbf{R}) \cdot \mathbf{n}) \quad (13)$$

86 This approach eliminates the finite-difference parameter Δ and its associated numerical errors,
87 while simultaneously providing forces \mathbf{F} and energy E in a single backward pass. The imple-
88 mentation reduces the computational cost per iteration from two force evaluations (standard
89 finite-difference) to one combined force + HVP evaluation.

90 Each dimer iteration consists of two phases:

91 **Rotation phase:** Minimize the curvature $\kappa = \mathbf{n}^T \mathbf{H} \mathbf{n}$ to align \mathbf{n} with the lowest-curvature
92 mode. The rotational force is:

$$\mathbf{F}_{\text{rot}} = (\mathbf{I} - \mathbf{n}\mathbf{n}^T) \mathbf{H}\mathbf{n} \quad (14)$$

93 The orientation is updated via gradient descent:

$$\mathbf{n}_{\text{new}} = \text{normalize}(\mathbf{n} - \alpha_{\text{rot}} \mathbf{F}_{\text{rot}}) \quad (15)$$

94 where $\alpha_{\text{rot}} = 0.5$ is the rotation step factor. Rotation converges when $\max |\mathbf{F}_{\text{rot}}| < 10^{-3} \text{ Eh/\AA}$ and
95 $\text{RMS}(\mathbf{F}_{\text{rot}}) < 5 \times 10^{-4} \text{ Eh/\AA}$.

96 **Translation phase:** Move along the effective force direction. The perpendicular force
97 component is:

$$\mathbf{F}_{\perp} = \mathbf{F} - (\mathbf{F} \cdot \mathbf{n})\mathbf{n} \quad (16)$$

98 When $\kappa < 0$ (indicating the mode is uphill), we perform a parallel flip to accelerate escape
99 from the convex region:

$$\mathbf{F}_{\text{trans}} = \begin{cases} \mathbf{F}_{\perp} - (\mathbf{F} \cdot \mathbf{n})\mathbf{n} & \text{if } \kappa < 0 \\ \mathbf{F}_{\perp} & \text{otherwise} \end{cases} \quad (17)$$

100 The position is updated with trust-region control:

$$\mathbf{R}_{\text{new}} = \mathbf{R} + \min(\alpha\mathbf{F}_{\text{trans}}, R_{\text{max}}) \quad (18)$$

101 where the step scaling α is adapted based on whether the step reaches the trust-radius boundary
102 ($R_{\text{max}} = 0.15 \text{ \AA}$).

Table S1: Summary of Current and Planned Functionalities in MAPLE

Category	Method	Description
Geometry Optimization	L-BFGS ^a	Batched GPU-accelerated quasi-Newton optimizer with dynamic batch shrinking
	RFO	Rational function optimization with adaptive trust-region control
	SD / CG	Steepest descent and conjugate gradient methods
	DIIS	Direct inversion in iterative subspace acceleration
Transition State Search	CI-NEB	Climbing-image nudged elastic band with IDPP interpolation and dynamic springs
	DS-P-RFO ^a	Dual-shift partitioned RFO with independent up-hill/downhill trust-region control
	Dimer ^a	Minimum-mode following with analytic Hessian-vector products via autodiff
	GSM	Growing string method with adaptive path construction and CI refinement
Intrinsic Reaction Coordinate (IRC)	AutoNEB	Automated multi-step NEB for complex reaction pathway exploration
	GS	Gonzalez-Schlegel method with mass-weighted coordinates
	HPC	Hessian predictor-corrector
	EulerPC	Euler predictor-corrector
Frequency & Thermochemistry	LQA	Local quadratic approximation
	Hessian Analysis	Analytic (autodiff) or numerical (finite-difference) mass-weighted Hessian
PES Scan	Thermochemistry	ZPE, enthalpy, entropy, Gibbs free energy; RRHO and quasi-RRHO ^b treatments
	Relaxed / Rigid	1D, 2D, and 3D scans along bond, angle, or dihedral coordinates
Single Point	Energy / Forces	Single-structure and batch trajectory energy and force evaluation
Molecular Dynamics	NVE	Velocity Verlet integrator for microcanonical ensemble
	NVT / NPT	Canonical and isothermal-isobaric ensembles with thermostats
Corrections & Solvation	DFT-D4	Dispersion correction with PyTorch automatic differentiation
	GBSA	Generalized Born surface area implicit solvation (OBC-II model)
	Explicit Solvent	Built-in library of pre-equilibrated explicit solvent environments

^a Algorithm described in the Methods section.

104 **Benchmark Dataset**

105 The 100 reactions used for benchmarking were drawn by uniform random sampling (seed = 42)
106 from the full Transition1x dataset⁵ (10,073 reactions), with no additional filtering criteria ap-
107 plied. All cases are unimolecular rearrangements of neutral, closed-shell singlet molecules
108 containing C, H, N, and O, spanning 43 unique molecular formulas with 6–19 atoms per
109 molecule (mean 13.7, median 14). This composition is representative of the full Transition1x
110 dataset, which was constructed by identifying spontaneous bond-breaking/forming events dur-
111 ing molecular dynamics simulations of small organic molecules. A detailed breakdown of the
112 elemental composition is provided in Table S2.

113 For the 100-reaction benchmark, the reactant and product structures were taken directly
114 from the Transition1x dataset without further optimization; the same endpoint geometries were
115 used for both DFT (ORCA) and MLFF (MAPLE) CI-NEB calculations.

Table S2: Composition of the 100-Reaction Benchmark Test Set

Elemental Composition	Count	Percentage
C, H, O	42	42%
C, H, N, O	33	33%
C, H, N	20	20%
C, H (hydrocarbons)	4	4%
C, N	1	1%
Total	100	100%

All reactions are unimolecular rearrangements of neutral, closed-shell singlet molecules. Molecular sizes range from 6 to 19 atoms (mean 13.7, median 14), with 3–6 carbon atoms per molecule. The 100 reactions span 43 unique molecular formulas and were drawn by uniform random sampling (seed=42) from the Transition1x dataset (10,073 reactions).

116 **Transition State Optimization**

Table S3: TS Optimization Performance Starting from MLFF-Derived HEI Structures, updated to include MACE-POL-S and MACE-POL-L.

Potentials ^a	RMSD ^b (Å)	Δ Iter ^c
ANI-1xnr (anal.)	0.61 ± 0.19	21.9
ANI-1xnr (num.)	0.61 ± 0.19	20.2
AIMNet2-NSE (anal.)	0.52 ± 0.14	45.1
AIMNet2-NSE (num.)	0.46 ± 0.14	26.2
UMA (num.)	0.13 ± 0.14	2.8
MACE-POL-S (anal.)	0.27 ± 0.27	6.6
MACE-POL-S (num.)	0.28 ± 0.28	4.4
MACE-POL-L (anal.)	0.18 ± 0.31	3.6
MACE-POL-L (num.)	0.20 ± 0.33	3.6

^aFrequency calculation methods: anal. = analytic frequencies via automatic differentiation; num. = numerical frequencies via finite differences. ^bGeometric deviation from DFT-optimized TS structures when TS optimization is restarted from MLFF-derived highest-energy image (HEI) structures. ^cIteration difference relative to DFT (MLFF – DFT).

Table S4: Decomposition of imaginary-frequency count mismatches relative to DFT reference TS calculations.

Potentials ^a	Acc. ^b	$ \Delta\nu_{\min} ^c$	N_+^d	Extra $-100 < \nu < 0^e$	Extra $\nu \leq -100^f$	N_-^g
ANI-1xnr (anal.)	6.2%	291 ± 265	90	48 (22%)	172 (78%)	0
ANI-1xnr (num.)	6.2%	292 ± 265	90	48 (22%)	171 (78%)	0
AIMNet2-NSE (anal.)	31.2%	529 ± 600	62	25 (22%)	89 (78%)	4
AIMNet2-NSE (num.)	32.3%	521 ± 567	62	25 (22%)	90 (78%)	3
UMA (num.)	86.5%	75 ± 109	11	7 (64%)	4 (36%)	2
MACE-POL-S (anal.)	47.9%	182 ± 187	46	27 (40%)	40 (60%)	4
MACE-POL-S (num.)	47.9%	182 ± 187	46	27 (40%)	40 (60%)	4
MACE-POL-L (anal.)	85.4%	54 ± 56	13	9 (69%)	4 (31%)	1
MACE-POL-L (num.)	86.5%	55 ± 55	12	9 (69%)	4 (31%)	1

^aSame methods as in the Hessian-quality benchmark. ^bPercentage of cases with the correct number of imaginary frequencies relative to the DFT reference. ^cMean absolute difference in the most negative imaginary frequency, computed only for cases where both the DFT and MLFF spectra contain at least one imaginary mode; the corresponding sample size is $n = 95$ for ANI-1xnr, AIMNet2-NSE, UMA, and MACE-POL-L, and $n = 93$ for MACE-POL-S. Units are cm^{-1} . ^dNumber of cases in which the MLFF predicts more imaginary frequencies than DFT. ^{e,f}When $N_{\text{ML}} > N_{\text{DFT}}$, the additional imaginary modes are defined as the shallowest $N_{\text{ML}} - N_{\text{DFT}}$ modes in the MLFF spectrum after sorting by $|\nu|$. Entries report the total number of such extra modes, with percentages relative to the total number of extra imaginary modes for that method. ^gNumber of cases in which the MLFF predicts fewer imaginary frequencies than DFT.

117 Kemp Eliminase Simulation

118 Protein Sequence and Structure

119 The amino acid sequence was extracted directly from the PDB structure (PDB ID: 5D2W)
120 and is provided below in FASTA format:

121 MALAKRIIAALIVKDGRVVKGSNFENLRDSGDPVELGKFYSEIGIDELSFWDITAS
122 VEKRKTMLELVEKVAEQIDIPFTVGGGIHDFETASELILRGADKVEINTAAVENPSLIT
123 QIAQTFGSQAVVVYIAAKRVDGEFMVFTYSGKKNTGILLRDWVVEVEKRGAGEIVL
124 GSIDRLGTKSGYDTEMIRFVRPLTTLPIAHGGAGKMEHFLEAFLAGADAANKANSVLH
125 FREIDVRELKEYLKKHGVNVRLEGLGSLEHHHHHH

Table S5: Experimental and Computational Barrier Heights for Kemp Eliminase Variants

Design	k_{cat} (s ⁻¹)	ΔG^\ddagger (kcal/mol)			Mutations ^a
		Obs. ^b	EVB ^c	MLFF ^c	
KE07	0.018 ± 0.001	20.1	19.5	17.9	WT
R6 3/7F	0.60 ± 0.070	17.8	16.1	17.8	I7D, K19E, K146T, G202R, N224D
R7 2/5B	1.20 ± 0.08	17.6	18.0	15.5	I7D, F77I, G202R, N224D
R7 1/3H	0.76 ± 0.03	17.8	17.0	17.0	I7D, V12L, F77I, H84Y, G202R, M207T, N222D, F229S

^aMutations relative to KE07 wild-type. ^bObserved activation barrier from experimental k_{cat} measurements. ^cCalculated activation barrier from computational method.

Table S6: Flexible Residue Definitions for Kemp Eliminase Simulations

Design	Flexible Residues ^a	Count
KE07 (WT)	9, 10, 11, 48, 49, 50, 51, 52, 80, 81, 101, 103, 128, 130, 144, 169, 171, 176, 201, 222, 224	21
R6 3/7F	9, 10, 11, 48, 49, 50, 51, 52, 80, 81, 101, 103, 128, 130, 144, 169, 171, 176, 201, 222, 224	21
R7 1/3H	9, 10, 11, 48, 49, 50, 51, 52, 80, 81, 101, 103, 128, 130, 144, 169, 171, 176, 201, 222, 224	21
R7 2/5B	9, 10, 11, 48, 49, 50, 51, 52, 80, 81, 101, 103, 128, 130, 144, 169, 171, 176, 201, 222	20

^aResidues within 8.0 Å of ligand center were defined as flexible and allowed to optimize during simulations. All other residues were kept fixed.




Figure S1: Overall structure of the wild-type KE07 Kemp Eliminate. Residues corresponding to mutation sites are highlighted in pink for reference, while the underlying structure remains the wild-type conformation.

126 **Diels-Alderase Simulation**

127 **Protein Sequence and Structure**

128 Protein structure of MaDAs are obtained from previous studies.^{6,7}

129 Protein sequence for MaDA-1: THEAFLECLTTRIPSNSTFTPQSIIYTPDNPSYSTILDS
130 TTQNPRFLSSSTRNPFAIITPLHASHIQAALYCSQKHGEQMRIRSGGDYEGLSYQSSVP
131 FFILDLRNLSSISIDAKSKSAWVQAGATIGELYYGIAKTSLNLSFPGGVAHTIGVGGQL
132 GGGGYGYSTRKYGLASDNVIDAQLIDARGRILDRKTMGEDLFWAIRGGGAGSFGIVL
133 AWKIRLVNTPSTVTIFEAVRSWENNTTKKFIRRYQRRASKTDKDLTIFVGFRTTSSTD
134 EEGNERISILTIVSATFHGSKDRLLQLVQKEFPDLGLVSEECTEMSWVRSIIHFNLFGDE
135 VPLEVLLNRTLNFEMKAFKLRSDYVQKPIPDDVLEKLLSKLYDEETGXGYIEFFPYGG
136 KMSKISESEIPFPYRAGNLYNLRYMVSWKDDGNITRTNMHLSWIKDAYDYMTPYVS
137 KDPRGAYLNFRDLDIGVNVNESDYDYVAKASVWGTTYFRNNFYRLVDIKTIVDPTN
138 FFKYEQSIPPLPPL

139 Protein sequence for MaDA-3: HESFLECLTTRISKSNSTSTPESIIYTKDNPSYSTILNS
140 TSLNPRFFPSSARYPLLIVTPLHASHVQATVHCAKKHGIQIRIRSGGDYEGLSYMSNVT
141 FAIVDLRNLSSIDVDVKKRKS AWVQSGATLGELHYWIAKKSQNLA FPGVVGHTVGIGG
142 MLGAGGGYGYSSRKYGLSADNILDAQLIDVRGRILNRKSMGEDLFWAIRGGGAGSFGI
143 VLAWKVRLVDVPSKVTVFTAIRDWDNNA TKKFIHRYQRRIAKVDKDLTIIVRFLTAS
144 ITDEKGSKKIQISTFITATYHGSQDRLLSLMEKEFPPELGLLAKECAEGAWVQSILYFNL
145 LTNSKSLDVLLNRTLNF EWRAF KIKSDYLKKPIPQVLENLLVKLYEEDIGETFVEFFP
146 YGGKLDEISESEIPCPHRAGNLYNLRYMVLWKEGQNATAVNKHL SWIRRAYNYMTP
147 YVSKNPRGAFLNFRDLDIGTNP NENEINGAYNYIKQASNWGTKYFKNNFYKLIYVKT
148 IVDPTNFFTYEQSIPSLPH




Figure S2: Overall structures of (a) MaDA-1 and (b) MaDA-3 showing the covalent linkage between residue H116 and the flavin adenine dinucleotide (FAD).




Figure S3: Two-dimensional potential energy surface (PES) analysis of the Diels-Alder reaction under different environments. (a) Definition of the two forming C–C bond distances, d_1 and d_2 , used as collective variables for the PES scans, illustrated on a representative reactive geometry. (b) Two-dimensional PESs projected onto the two forming C–C bond distances d_1 and d_2 (in Å) for the competing exo and endo pathways in vacuum, water, and ether. Color contours represent relative energies (kcal mol⁻¹), and red lines denote the minimum-energy paths determined using a minimum-energy path search algorithm.

Table S7: Flexible Residue Definitions for Diels-Alder Reaction Simulations

System	Flexible Residues ^a	Count
mada1	44, 45, 46, 47, 89, 91, 148, 149, 150, 151, 152, 153, 154, 159, 163, 164, 165, 166, 167, 168, 169, 174, 251, 263, 264, 265, 266, 269, 289, 290, 291, 326, 327, 328, 329, 330, 331, 332, 333, 342, 344, 345, 346, 347, 348, 349, 350, 351, 352, 353, 354, 355, 356, 357, 382, 383, 384, 385, 387, 388, 389, 390, 391, 414, 415, 416, 417, 418, 419, 420, 421, 422, 458, 459	74
mada3 ^b	42, 43, 44, 45, 46, 88, 89, 90, 91, 96, 130, 148, 149, 150, 151, 152, 153, 154, 165, 166, 167, 233, 234, 264, 265, 266, 267, 269, 287, 288, 289, 290, 325, 326, 327, 328, 329, 330, 331, 332, 333, 347, 348, 349, 350, 351, 352, 353, 354, 355, 382, 383, 384, 385, 386, 387, 388, 389, 390, 416, 417, 418, 419, 420, 421, 458, 459	67

^aResidues within 8.0 Å of ligand center were defined as flexible and allowed to optimize during simulations. All other protein framework residues were kept fixed. ^bThe same flexible residue set was used for mada3 and all its variants.

149 **Minimum Energy Path Determination**

150 The minimum energy path (MEP) connecting reactant and product configurations on the
151 2D free energy surface was determined using a Dijkstra-based algorithm combined with path
152 smoothing and uniform resampling. The free energy surface $F(x, y)$ was represented on a
153 regular 2D grid, with off-grid values obtained via bilinear interpolation.

154 The classic Dijkstra shortest-path algorithm⁸ was adapted to minimize the maximum en-
155 ergy barrier rather than total path length. Starting from the reactant minimum, the algorithm ex-
156 plores neighboring grid points using 8-connectivity (including diagonal moves) and maintains
157 for each point the lowest barrier B_{\max} encountered along any path from the start. The search
158 terminates upon reaching the product minimum, and the optimal path is reconstructed via back-
159 tracking. This approach guarantees finding the globally optimal minimum-barrier pathway on
160 the discrete grid with $O(N \log N)$ complexity, where N is the number of grid points.

161 The raw Dijkstra path exhibits staircase artifacts due to grid discretization. We applied
162 cubic B-spline smoothing ($k = 3$) to obtain a continuously differentiable path, parameterized
163 by cumulative arc length. The smoothed path was then uniformly resampled to N_{path} points
164 distributed with equal arc-length spacing:

$$s_i = \frac{i - 1}{N_{\text{path}} - 1} \cdot L_{\text{total}}, \quad i = 1, \dots, N_{\text{path}} \quad (19)$$

165 where L_{total} is the total path length. This uniform spacing ensures consistent resolution along the
166 entire reaction coordinate. Finally, the transition state location was refined by local grid search
167 in a small neighborhood around the maximum-energy point to accurately capture the barrier
168 height. For all calculations, $N_{\text{path}} = 15$ images were used, providing sufficient resolution for
169 subsequent nudged elastic band optimizations while maintaining computational efficiency.

170 The forward activation barrier $\Delta F_{\text{fwd}}^\ddagger$ and reverse barrier $\Delta F_{\text{rev}}^\ddagger$ were computed as the energy
171 difference between the transition state and reactant/product minima, respectively. The reaction
172 free energy is $\Delta F_{\text{rxn}} = F_{\text{product}} - F_{\text{reactant}}$.

173 DNA Photolyase Simulation

174 The amino acid sequence and DNA sequence were extracted directly from the PDB struc-
175 ture (PDB ID: 3CVU) and are provided below in FASTA format:

176 DNA Chain A: ACAGCGGTTGCAGGT

177 DNA Chain B: TACCTGCAACCGCTG

178 Protein sequence: MASWSHPQFEKGASTSLYKKAGLMDSQRSTLVHWFRKGLRLH
179 DNPALSHIFTAANAAPGRYFVRPIFILDPGILDWMQVGANRWRFLQQTLEDLDNQLR
180 KLNSRLFVVRGKPAEVFPRIFKSWRVEMLTFETDIEPYSVTRDAAVQKLAKAEGVRV
181 ETHCSHTIYNPELVIAKNLGKAPITYQKFLGIVEQLKVPKVLGVPEKLNMPPTPKDE
182 VEQKDSAAAYDCPTMKQLVKRPEELGPNKFPGGETEALRRMEESLKDEIWVARFEKP
183 NTAPNSLEPSTTVLSPYLKFGCLSARLFNQKLKEIIRKQPKHSQPPVSLIGQLMWREFY
184 YTVAAAEPNFDRLGNVYCMQIPWQEHDPHLEAWTHGRTGYPFIDAIMRQLRQEG
185 WIHHLARHAVACFLTRGDLWISWEEGQRVFEQLLLDQDWALNAGNWMWLSASAFF
186 HQYFRVYSPVAFGKKTDPQGHIYRKYVPELSKYPAGCIYEPWKASLVDQRAYGCVL
187 GTDYPHRIVKHEVVHKENIKRMGAAYKVNREVRTGKEEESSEFEEKSET

188 In the simulations, residues 288, 391, 394, 421, 423, 424, 425, 428, 429, 532, 8, and 9
189 were defined as flexible and allowed to optimize, while all DNA nucleotides except the pyrim-
190 idine(64)pyrimidone (64)-PP lesion site were kept fixed.




Figure S4: Overall structures of DNA Photolyase with DNA double strands.

191 **Multi-Step Reaction Pathway Construction**

192 The three-step repair mechanism (PCET, hydroxyl transfer, and C–C bond cleavage) was
193 modeled as sequential elementary reactions, each constructed and optimized individually. The
194 general workflow for each step is as follows:

195 (i) Endpoint definition. Reactant and product structures for each elementary step were
196 defined based on the established mechanism from QM/MM studies⁹. For the first step, the
197 globally optimized enzyme-substrate complex served as the reactant; for subsequent steps, the
198 optimized product of the preceding step was used. (ii) CI-NEB pathway optimization. CI-
199 NEB calculations were performed using 10 intermediate images with the settings described
200 above. When the optimization did not converge within the maximum number of iterations, the
201 last set of images was extracted and used as the starting point for continued optimization with
202 reset L-BFGS history. (iii) Pathway inspection and refinement. After each CI-NEB run, the
203 converged pathway was visually inspected. If individual images deviated significantly from
204 the expected minimum energy path, for example, exhibiting unphysical bond lengths, steric
205 clashes, or discontinuous geometric changes relative to neighboring images, these problematic
206 images were identified and removed. The remaining images were then used to reinitialize a
207 new CI-NEB calculation with interpolated replacements. This iterative refinement was repeated
208 until a smooth, physically reasonable pathway was obtained. (iv) Energy profile assembly. The
209 converged pathways from all three steps were concatenated to produce the overall reaction
210 energy profile, with the endpoint energies of consecutive steps matched to ensure continuity.

211 This workflow requires manual intervention and chemical judgment at several stages, re-
212 flecting the current state of the art for modeling complex multi-step enzymatic mechanisms.
213 Automating multi-step reaction discovery and pathway validation is an active development di-
214 rection for MAPLE.

215 **Enzymatic Systems Calculation Performance**

Table S8: Wall-clock time for a single L-BFGS optimization step on the three enzymatic systems using the UMA model. All calculations were performed on a single NVIDIA A100 80 GB GPU with an AMD EPYC 7742 CPU.

System	Num. of Atoms	GPU Memory (MB)	Wall Time Per Round (s)
Kemp Eliminate (R7-1)	3,936	15,309	145.0
Diels–Alderase (MaDA1)	8,351	18,696	170.8
DNA Photolyase	9,240	18,943	175.8

216 **Computational Details**

217 **Climbing-image Nudged Elastic Band Calculation**

218 Ab initio climbing-image nudged elastic band (CI-NEB) calculations were performed using
219 the ORCA-6.0 program package. All NEB calculations were carried out at the ω B97M-V/def2-
220 TZVPD level of theory.^{10,11} The reactant and product structures were provided as Cartesian
221 coordinates, and the reaction pathway was discretized into 10 intermediate images connecting
222 the reactant and product minima. Initial NEB pathways were generated by linear interpolation
223 of Cartesian coordinates optimized by the Image-Dependent Pair Potential (IDPP) method¹²
224 to produce physically reasonable geometries, particularly for reactions involving significant
225 bond rearrangements. Geometry optimizations along the reaction pathway were performed
226 using the CI-NEB algorithm to locate the minimum-energy path and corresponding transition-
227 state region. All calculations employed a closed-shell singlet electronic configuration. All
228 optimizations were performed using the default convergence criteria implemented in ORCA.

229 MAPLE CI-NEB calculations were performed using a fixed number of intermediate images
230 (10, excluding the end points) and optimized with a limited-memory BFGS (L-BFGS) algo-
231 rithm. Dynamic spring constants were employed, with values ranging from 0.03 to 0.3 Hartree \AA^{-2}
232 and a decay factor of 0.5, to improve path resolution near the saddle-point region. Initial reac-
233 tion pathways were generated using the image-dependent pair potential (IDPP) method. NEB
234 optimizations were converged based on projected forces, with thresholds of 9.5×10^{-3} and
235 5.0×10^{-3} Hartree \AA^{-1} for the maximum and root-mean-square projected forces, respectively.
236 A maximum of 500 optimization steps was allowed, using an L-BFGS history size of 20 and
237 an initial step length of 2×10^{-2} \AA . When CI-NEB refinement was applied, stringent force cri-
238 teria were used together with a reduced initial step size to accurately locate the highest-energy
239 image.

240 **Frequency Analysis**

241 Frequency analyses were performed in ORCA using numerical differentiation, with all
242 other settings kept at their default values.

243 In MAPLE, harmonic vibrational frequencies were computed from the Cartesian Hessian
 244 expressed in Hartree \AA^{-2} . Prior to diagonalization, the six external degrees of freedom cor-
 245 responding to overall translations and rotations were removed by projection in mass-weighted
 246 coordinates. The projected Hessian was constructed as

$$\mathbf{H}_{\text{proj}} = \mathbf{P}\mathbf{H}\mathbf{P}, \quad \mathbf{P} = \mathbf{I} - \mathbf{Q}\mathbf{Q}^{\text{T}}, \quad (20)$$

247 where \mathbf{Q} denotes the orthonormalized mass-weighted translationrotation basis.

248 The projected Hessian was then mass-weighted and diagonalized to obtain eigenvalues
 249 $\{\lambda_k\}$, which were converted to harmonic vibrational frequencies according to

$$\tilde{\nu}_k = \text{sgn}(\lambda_k) \sqrt{|\lambda_k|} C, \quad (21)$$

250 with $C = 2721.1383$ converting eigenvalues in (Hartree \AA^{-2})/amu to frequencies in cm^{-1} .

251 Normal modes were normalized in the mass-weighted metric following the ORCA convention.

252 The quality of Hessians was further evaluated by comparing the normal modes associated
 253 with the lowest imaginary frequency against DFT references. Two complementary metrics
 254 were used: the angular deviation and the mass-weighted mode overlap.

255 The mass-weighted overlap between an MLFF normal mode \mathbf{Q}_{ML} and the corresponding
 256 DFT mode \mathbf{Q}_{DFT} was defined as

$$\text{Overlap} = \frac{|\mathbf{Q}_{\text{DFT}}^{\text{T}} \mathbf{M} \mathbf{Q}_{\text{ML}}|}{\sqrt{(\mathbf{Q}_{\text{DFT}}^{\text{T}} \mathbf{M} \mathbf{Q}_{\text{DFT}}) (\mathbf{Q}_{\text{ML}}^{\text{T}} \mathbf{M} \mathbf{Q}_{\text{ML}})}}, \quad (22)$$

257 where \mathbf{M} is the diagonal mass matrix and \mathbf{Q} denotes the eigenvector of the mass-weighted Hes-
 258 sian matrix corresponding to a vibrational normal mode. The absolute value in the numerator is
 259 taken to eliminate the arbitrary sign of the normal-mode eigenvectors, ensuring that the overlap
 260 reflects physical similarity between modes.

261 The corresponding angular deviation was computed as

$$\theta = \arccos(\text{Overlap}), \quad (23)$$

262 where θ is reported in degrees.

263 **Transition State Optimization**

264 Transition-state optimizations in ORCA were performed starting from the converged CI-
265 NEB highest-energy image. Numerical Hessians were first evaluated at the ω B97M-V/def2-
266 TZVPD level of theory and the numerical Hessian was read into ORCA as the initial Hessian
267 guess. Subsequent transition-state optimizations were performed using ORCA's default set-
268 tings.

269 Transition-state optimizations in MAPLE were carried out using a trust-regionbased parti-
270 tioned rational function optimization (P-RFO) algorithm. A dual-shift formulation with mode
271 following was employed to explicitly locate first-order saddle points, where the reaction coordi-
272 nate was tracked by maximizing overlap between successive eigenmodes of the mass-weighted
273 Hessian. Hessians were obtained directly from the underlying MLFF calculators and used to
274 construct the quadratic model at each optimization step. The trust-region radii were dynam-
275 ically adjusted based on the agreement between predicted and actual energy changes. The
276 optimization was terminated when both force- and displacement-based convergence criteria
277 were satisfied, with thresholds defined on the maximum and root-mean-square values. Un-
278 less otherwise stated, all remaining algorithmic parameters were kept at their default MAPLE
279 settings.

280 **Explicit Solvent Potential Energy Surface Scan**

281 Reactant structures were taken from previous reports⁶ and initially optimized in MAPLE
282 using the L-BFGS algorithm with default convergence criteria. Starting from the optimized
283 geometry, the two forming CC bond distances were manually elongated to 4.0 Å and subse-
284 quently constrained during geometry optimization to generate the initial scanning structure.
285 This structure was directly used for gas-phase potential energy surface (PES) scans.

286 To best account for the solvent effect, explicit solvent models were employed. Solvent
287 boxes were constructed using the GAFF¹³ with ABCG2 charges.¹⁴ Each system was first sub-
288 jected to energy minimization followed by a 1 ns NVT molecular dynamics simulation to ob-

tain equilibrated solvent configurations. All solvent molecules within a $35 \times 35 \times 35 \text{ \AA}^3$ cube were extracted to form an equilibrium solvent box for the solvent. To accelerate the solvation process, MAPLE provides a library of nearly one hundred such pre-equilibrated solvents representing diverse environments.

Explicit solvent models were built using the MAPLE solvation builder, where solvent molecules were placed within a 10 \AA radius from the geometric center of the reactant. Solvent molecules located beyond 8 \AA were kept fixed to mimic the bulk environment, while the reactant and solvent molecules within 8 \AA were allowed to relax. Solvent molecules exhibiting atomic clashes within 1.5 \AA of the reactant were removed to eliminate any unrealistic forces that could compromise physically reasonable configurations. Constrained geometry optimizations were then performed to obtain the initial solvated structures in water and diethyl ether.

Two-dimensional PES scans were carried out by simultaneously varying the two CC bond distances from 4.0 to 1.5 \AA with a step size of 0.05 \AA , resulting in a total of $51 \times 51 = 2601$ sampling points.

Protein Structure Preparation

The raw PDB file was first processed using the `pdb4amber` module in AmberTools¹⁵ to repair common structural issues, including missing atoms, nonstandard residue naming, and to add hydrogen atoms consistent with standard protonation states. This procedure also ensured a standardized and simulation-ready PDB format.

Based on the repaired wild-type structure, site-specific residue mutations were introduced using the `tLeap` module in AmberTools, resulting in mutant variants. All mutations were applied while preserving the overall protein backbone conformation.

The corresponding small-molecule ligand present in the experimental protein–ligand complex was then incorporated into each system according to the crystallographic binding mode. The full protein–ligand complexes were subsequently imported into the MAPLE framework, where geometry optimizations of the entire systems were performed using the UMA machine-learning force field. Optimizations were carried out under superloose convergence criteria to efficiently relax the structures while maintaining the experimentally resolved binding geometry.

317 The optimized protein–ligand complexes obtained from these procedures were used as the
318 initial structures for all subsequent catalytic reaction calculations.

References

- (1) Liu, D. C.; Nocedal, J. On the limited memory BFGS method for large scale optimization. *Mathematical Programming* **1989**, *45*, 503–528.
- (2) Baker, J. An algorithm for the location of transition states. *Journal of Computational Chemistry* **1986**, *7*, 385–395.
- (3) Besalú, E.; Bofill, J. M. On the automatic restricted-step rational-function-optimization method. *Theoretical Chemistry Accounts* **1998**, *100*, 265–274.
- (4) Henkelman, G.; Jónsson, H. A dimer method for finding saddle points on high dimensional potential surfaces using only first derivatives. *The Journal of Chemical Physics* **1999**, *111*, 7010–7022.
- (5) Schreiner, M.; Bhowmik, A.; Vegge, T.; Busk, J.; Winther, O. Transition1x—a dataset for building generalizable reactive machine learning potentials. *Scientific Data* **2022**, *9*, 779.
- (6) Gao, L. et al. FADdependent enzymecatalysed intermolecular [4+2] cycloaddition in natural product biosynthesis. *Nat. Chem.* **2020**, *12*, 620–628.
- (7) Gao, L.; Zou, Y.; Liu, X.; Yang, J.; Du, X.; Wang, J.; Yu, X.; Fan, J.; Jiang, M.; Li, Y.; Houk, K. N.; Lei, X. Enzymatic control of endo and exostereoselective DielsAlder reactions with broad substrate scope. *Nat. Catal.* **2021**, *4*, 1059–1069.
- (8) Ensing, B.; Laio, A.; Parrinello, M.; Klein, M. L. A Recipe for the Computation of the Free Energy Barrier and the Lowest Free Energy Path of Concerted Reactions. *J. Phys. Chem. B* **2005**, *109*, 6676–6687.
- (9) Faraji, S.; Groenhof, G.; Dreuw, A. Combined QM/MM Investigation on the Light-Driven Electron-Induced Repair of the (6–4) Thymine Dimer Catalyzed by DNA Photolyase. *J. Phys. Chem. B* **2013**, *117*, 10071–10079.
- (10) Mardirossian, N.; Head-Gordon, M. ω B97M-V: A combinatorially optimized, range-separated hybrid, meta-GGA density functional with VV10 nonlocal correlation. *The Journal of Chemical Physics* **2016**, *144*, 214110.

- 345 (11) Weigend, F.; Ahlrichs, R. Balanced basis sets of split valence, triple zeta valence and
346 quadruple zeta valence quality for H to Rn: Design and assessment of accuracy. *Physical*
347 *Chemistry Chemical Physics* **2005**, *7*, 3297–3305.
- 348 (12) Smidstrup, S.; Pedersen, A.; Stokbro, K.; Jónsson, H. Improved Initial Guess for Mini-
349 mum Energy Path Calculations. *J. Chem. Phys.* **2014**, *140*, 214106.
- 350 (13) Wang, J.; Wolf, R. M.; Caldwell, J. W.; Kollman, P. A.; Case, D. A. Development and
351 Testing of a General Amber Force Field. *J. Comput. Chem.* **2004**, *25*, 1157–1174.
- 352 (14) He, X.; Man, V. H.; Yang, W.; Lee, T.-S.; Wang, J. ABCG2: A Milestone Charge Model
353 for Accurate Solvation Free Energy Calculation. *J. Chem. Theory Comput.* **2025**, *21*,
354 3032–3043.
- 355 (15) Case, D. A. et al. AmberTools. *J. Chem. Inf. Model.* **2023**, *63*, 6183–6191.