CATALYSTATE: ENZYME-CONDITIONED POSTERIOR TRANSPORT OVER GENERIC REACTION PRIORS

Anonymous authors
Paper under double-blind review

ABSTRACT

Recent transition-state generators have shown that reaction geometry can be learned as a transferable prior from reactant and product structures. Enzymatic reactions pose a different problem: the relevant transition region is constrained by a particular active site, while enzyme-specific transition-state structures are rarely available as supervision. We present CatalyState, which treats this setting as inference over a frozen reaction prior rather than as a new coordinate-generation problem. Given a generic transition-region proposal, CatalyState learns an endpoint-grounded enzyme–geometry preference and applies Catalytic Posterior Transport (CPT) to infer a constrained distribution over nearby geometries. This separates transferable reaction chemistry from enzyme-specific conditioning while explicitly enforcing mapped bond-edit, steric, and geometric constraints. Across enzyme- and reaction-disjoint evaluations, CatalyState distinguishes native catalytic environments from controlled context perturbations, improves local proposals from multiple distinct reaction priors, and transfers its learned preference to an independent semiempirical reaction-coordinate evaluator. Matched local-search controls further show that this preference supports within-system refinement without acting as a surrogate physical energy. Together, these results support adapting, rather than relearning, transferable reaction priors for enzyme-specific reaction geometry.

1 INTRODUCTION

A reaction model and an enzyme constrain different aspects of the same chemical event. The reaction specifies the bond changes that must occur; the enzyme determines which geometric realizations are compatible with a particular catalytic environment. Machine learning has increasingly treated transition-state search as a geometric prediction or generation problem. Quantum-chemical datasets provide reactant/product/transition-state supervision at scale (Grambow et al., 2020; Schreiner et al., 2022); direct predictors and learned second-order information accelerate initialization or refinement (Choi, 2023; Yuan et al., 2024); and diffusion, optimal-transport, and flow models generate transition states or reaction paths from reaction information (Duan et al., 2023; Kim et al., 2024; Duan et al., 2025; Hayashi et al., 2025; Luo et al., 2026; Tuo et al., 2026). Yet generalization to genuinely unseen chemistry remains difficult (Darouich et al., 2026), and these models are primarily reaction-conditioned rather than organized around the catalytic residues, steric restrictions, and local chemistry of a specific active site. Transferable reaction geometry is therefore a useful prior for enzyme-specific refinement, but not a complete enzyme-conditioned solution.

Enzyme modeling supplies complementary information. Recent work addresses function annotation, substrate specificity, reaction-conditioned retrieval, active-site identification, and biomolecular complex structure prediction (Yu et al., 2023; Kroll et al., 2023; Mikhael et al., 2024; Song et al., 2024; Wang et al., 2024; Cui et al., 2025; Stärk et al., 2022; Corso et al., 2023; Qiao et al., 2024; Lu et al., 2024; Abramson et al., 2024), while structure-based screening shows that catalytically favorable active-site conformations can add information beyond sequence similarity (Malatesta et al., 2024). Biochemical resources provide curated reactions, mapped transformations, contacts, and catalytic-site evidence (Chang et al., 2021; Bansal et al., 2022; Heid et al., 2023; Zhang et al., 2024; Ribeiro et al., 2018), and modern atom mapping makes the net graph transformation explicit (Schwaller et al., 2021a; Chen et al., 2024). What remains missing is direct supervision for the enzyme-specific transition region itself, which is transient and typically inferred through kinetics, isotope effects, or computation (Garcia-Viloca et al., 2004; Schramm, 2011). This creates an asymmetric learning

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

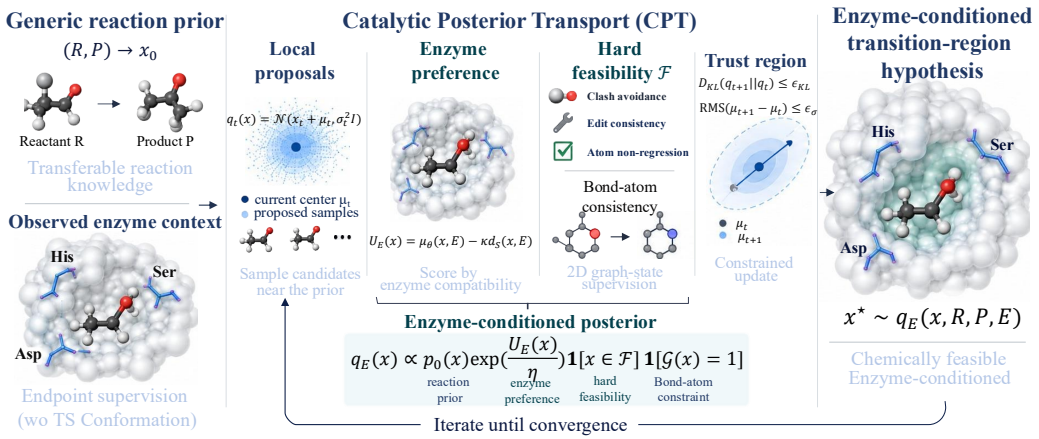


Figure 1: **CatalyState overview.** Endpoint complexes provide enzyme-context supervision for a local preference model. Given a transition-region proposal from a frozen reaction prior, CPT samples nearby geometries, removes candidates that violate mapped bond-edit or geometric constraints, and performs bounded preference-guided updates to return an enzyme-conditioned hypothesis.

problem: rich endpoint context must guide an unobserved local reaction geometry. Importantly, recognizing a plausible endpoint complex is not sufficient by itself. The useful signal must distinguish how nearby reaction geometries interact with the catalytic environment while preserving the bond changes already encoded by the reaction model. This is the regime in which endpoint supervision can act as a *local preference* rather than as a surrogate transition-state label. Throughout, we use *transition-region geometry* for the local geometric hypothesis refined by our method and reserve *transition state* for a physically validated saddle-point structure.

We address this problem with **CatalyState**. A frozen reaction model supplies a generic proposal, an endpoint-grounded critic supplies an enzyme–geometry preference, and Catalytic Posterior Transport (CPT) searches a bounded neighborhood for higher-utility configurations under explicit support constraints. The decomposition separates three roles: *preference* ranks nearby configurations, *hard support* removes inadmissible ones, and a *trust region* limits departure from the transferable reaction prior. Mapped bond edits identify the reacting subspace and local geometric constraints, while supported catalytic annotations provide additional enzyme context.

The main contribution is this problem decomposition: enzyme-specific transition-region refinement is treated as *local conditional specialization of a frozen reaction prior*, rather than as a new coordinate-generation problem. We establish the decomposition through three complementary tests. Controlled context perturbations ask whether endpoint supervision actually encodes enzyme information; matched refinement across multiple upstream priors asks whether that information transfers to nearby reaction geometries; and an independent semiempirical evaluator tests whether the learned ordering extends beyond the critic's own score. Together with matched local-search controls, these experiments characterize both the utility and the limits of endpoint-derived enzyme preferences.

2 Related Work

Reaction geometry and transition-state modeling. Classical transition-state search relies on saddle-point and path optimization, including nudged-elastic-band variants (Henkelman & Jónsson, 2000; Henkelman et al., 2000). The learned reaction-modeling toolkit includes quantum reaction datasets and atomistic potentials (Grambow et al., 2020; Schreiner et al., 2022; Schütt et al., 2017; Batzner et al., 2022). More task-specific methods perform direct transition-state prediction or learned-Hessian refinement (Choi, 2023; Yuan et al., 2024), and diffusion, optimal-transport, or flow formulations generate transition states and reaction paths (Duan et al., 2023; Kim et al., 2024; Duan et al., 2025; Hayashi et al., 2025; Luo et al., 2026; Tuo et al., 2026). OOD studies show that strong in-domain generation does not guarantee transfer to unseen chemistry (Darouich et al., 2026).

CatalyState is complementary: it keeps this transferable reaction knowledge frozen and specializes a generic proposal to a particular enzyme environment.

Reaction representations and biochemical resources. Mechanistic and graph-based representations provide the discrete reaction evidence needed to localize such refinement (Jorner et al., 2021). Atom-mapping methods recover atom correspondences and net bond rearrangements (Jaworski et al., 2019; Schwaller et al., 2021a; Chen et al., 2024); graph-edit models identify reaction centers (Jin et al., 2017; Coley et al., 2019); and learned reaction representations organize reaction outcomes, reaction space, or electron redistribution (Schwaller et al., 2019; 2021b; Joung et al., 2025). For enzymatic chemistry, BRENDA, Rhea, and EnzymeMap provide curated functional and reaction information (Chang et al., 2021; Bansal et al., 2022; Heid et al., 2023). CatalyState uses mapped bond edits to define the reacting subspace and its local geometric support.

Enzyme context and complex geometry. A separate line of work addresses protein function, enzyme–substrate compatibility, reaction-conditioned retrieval, and active-site prediction (Gligorijević et al., 2021; Yu et al., 2023; Kroll et al., 2023; Mikhael et al., 2024; Song et al., 2024; Cui et al., 2025; Wang et al., 2024); CSA, M-CSA, and BioLiP2 provide experimentally grounded catalytic-site and ligand-contact evidence (Porter et al., 2004; Ribeiro et al., 2018; Zhang et al., 2024). Protein–ligand methods further span docking, pocket-conditioned generation, and flexible-complex prediction (Stärk et al., 2022; Lu et al., 2022; Peng et al., 2022; Cao et al., 2024; Zhang et al., 2023b; Corso et al., 2023; Qiao et al., 2024; Lu et al., 2024; Abramson et al., 2024). Physical-validity studies such as PoseBusters (Buttenschoen et al., 2024) and structure-based enzyme screening (Malatesta et al., 2024) motivate explicit geometric checks and enzyme-aware conditioning. These approaches provide rich endpoint or binding context, but do not directly specify how that context should modify a frozen reaction-region proposal.

Constrained inference and computational enzymology. Posterior regularization, relative-entropy policy search, trust-region optimization, and the cross-entropy method provide generic tools for constrained inference and bounded distributional search (Ganchev et al., 2010; Peters et al., 2010; Schulman et al., 2015; Rubinstein & Kroese, 2004). CatalyState uses these ideas as inference machinery; its contribution is the decomposition of a frozen reaction prior, an endpoint-grounded enzyme preference, mapped bond-edit support, and bounded local transport. At the physical end, transition-state theory, isotope effects, and QM/MM connect enzyme catalysis to barrier-region structure and energetics (Garcia-Viloca et al., 2004; Schramm, 2011; Warshel et al., 2006; Senn & Thiel, 2009; van der Kamp & Mulholland, 2013; Zhang et al., 2023a). Together, these lines of work leave open the interface targeted here: injecting sparse enzyme-specific context into transferable reaction geometry without retraining the underlying reaction model.

3 Problem Formulation

Inputs and output. The input comprises a mapped reaction (R, P), enzyme context E, endpoint-coordinate information, and a frozen generic reaction proposal $x_0 = T(R, P)$. CatalyState returns a local reaction-region hypothesis, optionally together with additional diverse admissible hypotheses. The pocket frame is fixed during refinement, so the task is local specialization of an upstream proposal rather than de novo complex construction.

Local reference and admissibility. Let n be the number of movable reacting atoms. On their $3n$-dimensional coordinate subspace, we define

$$p_{0,b}(x) = \mathcal{N}(x_0, \sigma_b^2 I), \qquad \sigma_b = 0.30b/\sqrt{3}, \tag{1}$$

where b is the allowed reacting-atom RMS displacement; all other coordinates remain fixed. The admissible set $\mathcal{F}_b$ combines finite coordinates, the displacement cap, severe-clash exclusion, mapped edited-pair distance intervals, and edit/available-actor non-regression checks. These constraints define the local support within which enzyme conditioning is allowed to modify the upstream geometry.

A constrained reference posterior. For structural utility $U_E(x)$ and $\eta > 0$, consider

$$\max_{q:\,\mathrm{supp}(q)\subseteq\mathcal{F}_b} \left\{\mathbb{E}_q[U_E(x)] - \eta D_{\mathrm{KL}}(q\|p_{0,b})\right\}. \tag{2}$$

When the feasible prior mass and normalizer are finite and nonzero, the solution is the Gibbs tilt

$$q_\eta^\star(x) = \frac{1}{Z_\eta}\, p_{0,b}(x) \exp(U_E(x)/\eta)\, \mathbf{1}[x \in \mathcal{F}_b]. \tag{3}$$

This standard variational construction (Ganchev et al., 2010; Peters et al., 2010) makes the decomposition explicit: the frozen prior supplies transferable reaction geometry, U_E supplies enzyme-specific preference, and $\mathcal{F}_b$ defines admissible local support. Appendix B gives the corresponding support conditions.

Practical approximation. The implemented search maintains an untruncated Gaussian proposal $q_t = \mathcal{N}(x_0 + \mu_t, \sigma_b^2 I)$ and updates its mean using feasible samples. CPT-REPS performs a relative-entropy reweighting of the current feasible population, whereas CPT-CEM updates from feasible high-utility elites. Both are finite-horizon approximations to the same prior–preference–support decomposition, with explicit movement limits defined below.

4 CATALYSTATE

4.1 ENDPOINT-GROUNDED PREFERENCE AND REACTION EVIDENCE

A three-member E(3)-invariant ensemble scores reacting-atom–pocket pairs using reacting-atom identity, protein residue type, protein atom type, radial distance features, and an actor flag. Mapped reaction graphs identify edited atoms and edit types, while supported catalytic annotations provide additional local context. Mean/max pooling with local distance and actor statistics produces a scalar score; each ensemble member uses hidden width 64, 16 radial features, and an 8 Å interaction cutoff. Sample and protein identifiers are never provided to the critic, so the preference must be expressed through local chemical and geometric features rather than memorized system identity.

For ensemble member f_m, the endpoint ranking loss uses $\mathrm{softplus}(f_m(z^-) - f_m(z^+))$, averaged over members and available counterfactual families. Training contrasts native endpoints with local/far rigid displacements, pocket-chemistry permutations, actor reassignment, and supported role permutations, and additionally ranks local above far displacements. A generic reaction-prior frame is not used as a positive training example: the critic is learned from endpoint-grounded contrasts and only later applied to transition-region proposals. This separation is important because it prevents the upstream generator from supervising the very preference used to refine it. Matched foreign-pocket controls and a dual-critic variant are evaluated separately in Appendix G.

The search utility combines ensemble preference with an endpoint-actor residual:

$$U_E(x) = \underbrace{\bar{f}_\psi(x, E) - \kappa s_\psi(x, E)}_{\text{ensemble preference}} + \beta A_E(x), \qquad \kappa = 1, \quad \beta = 0.020. \tag{4}$$

Here s_ψ is the standard deviation across ensemble members, and A_E penalizes departure from the midpoint of endpoint reacting-atom–actor nearest distances where actor evidence is available. The disagreement term therefore makes uncertain proposals less attractive without interpreting ensemble variance as a calibrated physical uncertainty. Edit and clash constraints remain outside the utility, keeping preference and admissibility separate.

Bond-edit support. Mapped reactant/product graphs determine the edited atom pairs and their bond-order changes; the union of these atoms defines the reacting set. Hard support checks edited-pair geometry, bond-edit progress, severe clashes, and available actor contact under fixed tolerances. Appendix D specifies the exact thresholds.

4.2 FEASIBLE LOCAL POPULATION UPDATES

At each round, CPT draws 64 antithetic particles, $x_i = x_0 + \mu_t + \sigma_b \epsilon_i$, discards samples outside $\mathcal{F}_b$, and updates the proposal mean from the feasible population. If initial feasible support is sparse, a bounded preflight restoration phase shifts the mean using smooth edit/clash barriers while leaving the hard thresholds unchanged.

For CPT-REPS, feasible samples $\mathcal{I}_t$ receive

$$w_i(\eta_t) = \frac{\exp(U_E(x_i)/\eta_t)}{\sum_{j\in\mathcal{I}_t}\exp(U_E(x_j)/\eta_t)}, \qquad D_{\mathrm{KL}}(w\|\mathrm{Unif}(\mathcal{I}_t)) \le 0.50. \tag{5}$$

The temperature is solved per system and round. With the normalized weights held fixed during differentiation, REPS uses the pathwise direction

$$g_t = \nabla_{\mu_t}\sum_{i\in\mathcal{I}_t}\mathrm{sg}(w_i)U_E(x_i),$$

so the proposal moves toward particles preferred by the enzyme-conditioned utility while the relative-entropy constraint controls how concentrated the reweighting can become. CPT-CEM instead updates from high-utility feasible elites at fixed covariance. CEM serves as the practical optimizer; REPS is the information-constrained variant most directly aligned with Eq. 2. Fixed-temperature reweighting, uniform search, and projected gradient provide matched local-search controls.

Movement limits. For equal-covariance Gaussian proposals,

$$\frac{1}{n}D_{\mathrm{KL}}(q_{t+1}\|q_t) = \frac{\mathrm{RMS}(\mu_{t+1}-\mu_t)^2}{2\sigma_b^2} \le 0.25. \tag{6}$$

A separate projection imposes $\mathrm{RMS}(\mu_{t+1}) \le 0.80b$ relative to x_0, while every accepted particle must satisfy the full RMS cap b. These constraints distinguish per-step movement, accumulated mean movement, and realized particle displacement.

Operational convergence. We use a fixed refinement horizon for all matched CPT variants and record proposal movement and weighted empirical utility $\widehat{U}_t$ at each round. A trajectory is marked operationally converged when, for two consecutive active rounds,

$$\delta_\mu^{(t)} = \mathrm{RMS}(\mu_{t+1}-\mu_t) \le 10^{-4}\ \text{Å}, \qquad \delta_U^{(t)} = |\widehat{U}_t - \widehat{U}_{t-1}| \le 10^{-5}, \tag{7}$$

with sufficient feasible support and non-flat utility. This diagnostic characterizes finite-horizon stabilization; implementation details are given in Appendix C.

Output. After refinement, CPT samples a final proposal pool, reapplies all hard checks, and requires $U_E(x) \ge U_E(x_0)$. The highest-utility accepted particle is the primary output; up to eight diverse hypotheses can also be retained by reacting-atom RMS separation. If no particle satisfies both feasibility and the utility floor, the system is recorded as a non-selection rather than being silently repaired or counted as an improvement. This distinction is used directly in the feasible-selection rates reported below. We report solver identity explicitly because CEM and REPS implement different transport rules over the same critic, support, and movement limits.

5 Experimental Setup

Two supervision protocols. The endpoint corpus contains 4,236 deduplicated physical systems linked from EnzymeMap to BioLiP/RCSB holo evidence, yielding 1,992/95/863 accepted train/validation/test views. Enzyme groups use MMseqs2 clusters at 40% identity and 80% coverage; substrate groups use exact Bemis–Murcko scaffolds, with canonical-molecule identity for acyclic substrates. The finalized group-based split yields 287 enzyme-OOD, 274 substrate-OOD, and 302 double-OOD test views.

The local-refinement benchmark is a 171-system frozen paired-coordinate cache assembled for this study. Five connected-component folds of the enzyme–reaction graph have sizes 35/34/34/34/34; reaction clusters use EC level 2 together with the multiset of net edit types and bond-order changes. Each fold trains its own three-member critic on the remaining endpoint records, keeping the endpoint OOD benchmark and coordinate-level refinement benchmark as distinct evaluation protocols. Actor-specific controls use only systems with the required annotation. Full provenance and split construction are given in Appendix A.

Matched search protocol. React-OT, TS-DFM, and MolGEN provide distinct upstream reaction priors. We use three sampling seeds and four budgets $b \in \{0.05, 0.10, 0.25, 0.50\}$ Å, producing 2,052 system/seed/budget rows per prior. All local-search controls share the upstream structures, out-of-fold critic, hard eligibility checks, refinement horizon, and proposal budget; sampling seeds vary search randomness rather than critic training. Thus repeated search seeds probe stochastic refinement under a fixed learned preference, whereas the separate endpoint ablation table uses independent critic-training runs. We keep these two repetition axes distinct throughout the paper.

For selected proposal x, normalized learned gain is

$$G(x) = \frac{U_E(x) - U_E(x_0)}{\max\{|U_E(x_0) - U_E(x^{0.5}_{\text{decoy}})|, 10^{-4}\}}, \tag{8}$$

where the decoy rigidly displaces the movable reacting geometry by 0.5 Å away from its actor/pocket reference. When the denominator floor is inactive, a gain of one corresponds to the absolute native-to-decoy utility gap, making gains more comparable across systems with different score scales. Main search summaries select each method's highest-utility result over the four budgets, compute the median across systems for each seed, and then average across seeds; feasible-selection rates are reported separately. These summaries therefore measure best-of-budget search behavior, while the independent physical comparison below fixes $b = 0.25$ Å to remove budget selection from that evaluation.

Independent physical evaluator. We separately compare selected candidates with the interior maximum x_B of an endpoint-based restrained concerted GFN2-xTB scan (Bannwarth et al., 2019). For the mapped edited-pair set $\mathcal{E}$,

$$\text{MAE}_{\text{rxn}}(x) = \frac{1}{|\mathcal{E}|} \sum_{e \in \mathcal{E}} |d_e(x) - d_e(x_B)|. \tag{9}$$

The fixed-budget comparison uses $b = 0.25$ Å on matched usable systems within each prior. We report paired medians, system-bootstrap intervals, and strict lower-error fractions with denominators shown explicitly. A higher-coverage React-OT analysis additionally evaluates 194 protonation-specific profiles. The semiempirical calculations are used only as an external evaluator of local reaction-coordinate agreement.

6 Results

6.1 Does endpoint supervision contain enzyme-context information?

Endpoint supervision is consistently sensitive to enzyme context. Reassigning catalytic actors lowers the score relative to the native endpoint on 87.4–88.3% of enzyme-, substrate-, and double-OOD views, while fixed-geometry pocket-chemistry permutations do so on 95.5–98.5%. On role-supported subsets, role permutation lowers the score on 73.7–92.3% of views. Because geometry is held fixed in the pocket-chemistry intervention, these results show that the critic uses local enzyme identity and chemistry rather than geometry alone.

The broader OOD discrimination task is consistent with these interventions. Across the same three splits, the three-run full-critic AUPRC ranges from 80.02% to 86.85% (Table 1); removing auxiliary catalytic annotations or actor features lowers the corresponding point estimates, while geometry alone is substantially weaker. Together, these controls show that endpoint supervision contains a transferable enzyme-context signal.

6.2 Does local specialization transfer across reaction priors?

The learned preference transfers across all three upstream reaction priors. Table 2 compares matched local-search controls on React-OT, TS-DFM, and MolGEN. CPT-CEM gives the largest normalized learned-utility gain on each prior, with only a marginal advantage over REPS on React-OT and a clearer separation on TS-DFM and MolGEN. Fixed-temperature weighting remains competitive, indicating that the gain is not explained solely by adaptive temperature calibration. We therefore use CEM as the practical optimizer and retain REPS as the explicitly information-constrained variant.

Table 1: **Endpoint OOD and context controls.** Panel A reports mean ± s.d. over three training runs. Panel B reports the native-score win rate over controlled counterfactuals; pocket chemistry is permuted at fixed geometry.

	Enzyme OOD	Substrate OOD	Double OOD
A. Endpoint counterfactual AUPRC			
Full	83.12 ± 1.00	86.85 ± 1.48	80.02 ± 1.18
Without auxiliary catalytic annotations	81.57 ± 3.71	82.51 ± 6.25	76.79 ± 4.43
Without actor features	77.29 ± 1.09	84.85 ± 1.26	76.76 ± 0.31
Geometry only	62.59 ± 1.31	70.10 ± 0.65	64.99 ± 0.58
B. Native score > controlled context perturbation			
Actor reassignment	87.80%	88.32%	87.42%
Pocket-chemistry permutation	95.47%	98.54%	96.36%
Role permutation	92.31% ($n = 39$)	73.68% ($n = 38$)	75.76% ($n = 66$)

Table 2: **Matched local search across reaction priors.** Normalized learned-utility gain is shown with feasible-selection rate in parentheses, using each method's best utility over four budgets.

Search	React-OT	TS-DFM	MolGEN
Uniform	1.012 (96.7%)	0.621 (74.9%)	0.738 (78.9%)
Projected gradient	1.595 (93.0%)	0.945 (70.2%)	1.122 (77.2%)
Fixed temperature	2.437 (97.1%)	1.332 (73.5%)	1.497 (78.9%)
CPT-REPS	2.506 (97.1%)	1.317 (73.3%)	1.516 (78.9%)
CPT-CEM	2.512 (97.1%)	1.437 (74.5%)	1.674 (78.9%)

Feasible improvement depends on the upstream prior: CEM selects an admissible higher-utility candidate on 97.1% of React-OT systems and 74.5%/78.9% of TS-DFM/MolGEN systems. Uniform search attains a similar TS-DFM feasible-selection rate but substantially lower utility. Under matched wrong-pocket search, the correct enzyme context yields higher REPS gain on 94.7%, 72.5%, and 78.9% of React-OT, TS-DFM, and MolGEN systems, respectively. Thus transfer is not merely prior-agnostic score optimization: the refinement remains sensitive to the enzyme environment in which it is performed. The variation in feasible-selection rate is also informative. CatalyState does not erase differences between upstream generators; it can only exploit enzyme-compatible directions that remain inside the local support around each proposal. React-OT therefore offers more usable local mass under the declared constraints, whereas TS-DFM and MolGEN more often reach the support boundary before an improved candidate is found. This behavior is consistent with the intended modularity of the method: the reaction prior determines where refinement starts and which nearby geometries are reachable, while the enzyme preference determines how those reachable candidates are ranked.

6.3 Does learned improvement transfer to independent geometry?

The learned improvement partially transfers to an independent semiempirical reaction-coordinate proxy. At $b = 0.25$ Å, CPT-REPS lowers edited-pair MAE relative to the raw prior on 81/110 React-OT, 45/76 TS-DFM, and 54/78 MolGEN systems (Table 3). The median paired differences are -0.0246, -0.0112, and -0.0175 Å, respectively; the React-OT and MolGEN intervals exclude zero, while the TS-DFM interval reaches zero. The effect therefore extends beyond the learned score, but with different strength across priors. This distinction matters because learned-utility improvement alone could be produced by optimizing imperfections of the critic. Agreement with an evaluator that is neither used to train the critic nor used during search provides a separate test of whether the local ordering contains physically relevant geometric information.

The larger React-OT path analysis is directionally consistent: 137/194 protonation-specific profiles improve, and median reaction-coordinate error decreases from 0.5401 to 0.4750 Å (12.0%). In a separate restrained-relaxation comparison, median reaction-distance drift decreases from 0.0687 to 0.0363 Å across 166 complete pairs. By contrast, fixed-frame coordinate error changes from 0.8495 to 0.8659 Å, indicating that the improvement is localized to the reaction coordinate rather than to the global pose.

Table 3: **Independent fixed-budget reaction-coordinate evaluation.** At 0.25 Å, Δ is REPS minus reference MAE; negative favors REPS. Intervals are 95% system-bootstrap CIs.

Reference	n	REPS lower MAE	Median Δ	95% CI
Across upstream priors				
React-OT raw	110	73.6%	-0.0246	$[-0.0543, -0.0129]$
TS-DFM raw	76	59.2%	-0.0112	$[-0.0187, 0.0000]$
MolGEN raw	78	69.2%	-0.0175	$[-0.0336, -0.0027]$
React-OT solver controls				
CPT-CEM	110	31.8%	$+0.0126$	$[+0.0092, +0.0195]$
Uniform	109	45.9%	$+0.0018$	$[-0.0035, +0.0077]$
Projected gradient	109	22.9%	$+0.0593$	$[+0.0342, +0.0777]$

Matched local-search controls clarify what the transport constraint contributes. On React-OT, CEM lowers reaction-coordinate error relative to REPS by a median 0.0126 Å without a detectable path-normal difference. Projected gradient lowers edited-coordinate error by 0.0593 Å relative to REPS, whereas REPS has 0.0041 Å lower path-normal displacement (95% interval $[-0.00615, -0.00056]$). More aggressive direct optimization can therefore fit the measured coordinate more closely, while constrained transport better preserves the upstream geometric neighborhood.

6.4 What does the learned utility capture?

The critic behaves as a local ranking signal rather than a globally calibrated energy model. In a retrospective paired analysis that fixes prior, system, search seed, and protonation state, higher utility agrees with lower xTB path error on 67.5% of informative React-OT pairs, 59.3% of TS-DFM pairs, and 63.4% of MolGEN pairs; the corresponding system-bootstrap intervals are [63.3,71.6], [54.2,64.3], and [58.8,68.0]%. CPT therefore does not require absolute score comparability across unrelated enzymes: it requires a useful within-system ordering of nearby candidates. These analyses use 115/78/96 systems for React-OT/TS-DFM/MolGEN, respectively, and therefore need not coincide with the matched intersections used by the fixed-budget table. Together with the direct-search controls, the result identifies the learned preference as a transferable local ranking signal that can guide bounded refinement without serving as a physical energy model. The distinction between ranking and calibration is central to the method: CPT only compares nearby candidates within the same enzyme system, so useful local ordering is sufficient even when score magnitudes are not comparable across unrelated enzymes.

7 Discussion

The results suggest a broader modeling principle: when domain-specific supervision is scarce but a strong generic generator already exists, specialization need not require retraining the generator itself. Instead, domain-specific evidence can define a local preference over the generic prior, while explicit support and movement constraints limit how far adaptation departs from transferable structure. CatalyState instantiates this principle for enzymatic reaction geometry, using endpoint supervision to condition reaction priors that were learned without enzyme-specific transition-region targets. This modular view has a practical consequence: improvements in upstream reaction generation and improvements in enzyme-context modeling can be developed independently and composed at inference time. A better reaction prior expands the quality of the local neighborhood; a better enzyme preference improves how that neighborhood is ordered.

The solver comparisons clarify the role of the transport mechanism within this decomposition. CEM gives the best learned-utility performance among the tested solvers and also improves the matched React-OT reaction-coordinate result over REPS. REPS imposes a stronger relative-entropy bias toward the current population, while projected gradient can optimize the edited coordinate more aggressively at the cost of larger orthogonal displacement. The resulting accuracy–conservatism trade-off shows why the scientific contribution is better viewed as enzyme-conditioned local specialization than as a claim about one universally optimal optimizer. It also explains why hard support and trust-region structure are useful even when a direct gradient step can optimize one measured coordinate

more aggressively: the objective of specialization is to improve an enzyme-relevant direction without discarding all structure supplied by the transferable prior.

The present study targets local specialization once an upstream reaction geometry, mapped bond edits, and a pocket frame are available. The 171-system refinement benchmark is a coordinate-level testbed, and the GFN2-xTB analysis provides an external semiempirical proxy for reaction-coordinate agreement. Within this scope, enzyme context systematically improves transferable reaction proposals without requiring enzyme-specific transition-state coordinates as training targets. The current fixed-pocket setting isolates this question deliberately; allowing enzyme relaxation would couple preference learning with conformational response and would require stronger supervision or physical regularization. Extending the same decomposition to flexible pocket inference and higher-level energetic supervision would therefore test whether the local preference can remain useful as progressively more of the catalytic environment is allowed to move.

8 Conclusion

CatalyState reframes enzyme-specific reaction geometry as local specialization of a transferable reaction prior. By combining an endpoint-grounded enzyme preference with explicit geometric support and bounded transport, it introduces enzyme context without discarding reaction knowledge learned from broader chemical data. More broadly, scarce enzyme-specific transition-region supervision need not require reaction geometry to be learned from scratch: transferable chemical priors can instead be adapted locally to the catalytic environment.

AI use statement

Generative AI tools were used for limited supporting tasks, primarily literature organization, code debugging, and language and figure editing. They were not used to generate experimental data or to make scientific decisions. The authors designed and executed the study, verified the analyses and references, interpreted the results, and take responsibility for the final manuscript.

Reproducibility statement

Sections 3–5 define the task, method, splits, and evaluation. Appendices A–E document data provenance, implementation details, constraints, physical controls, and statistical units.

References

Josh Abramson, Jonas Adler, Jack Dunger, et al. Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature*, 630:493–500, 2024. doi: 10.1038/s41586-024-07487-w.

Christoph Bannwarth, Sebastian Ehlert, and Stefan Grimme. GFN2-xTB—an accurate and broadly parametrized self-consistent tight-binding quantum chemical method with multipole electrostatics and density-dependent dispersion contributions. *Journal of Chemical Theory and Computation*, 15 (3):1652–1671, 2019. doi: 10.1021/acs.jctc.8b01176.

Parit Bansal, Anne Morgat, Kristian B. Axelsen, Venkatesh Muthukrishnan, Elisabeth Coudert, Lucila Aimo, Nevila Hyka-Nouspikel, Elisabeth Gasteiger, Arnaud Kerhornou, Teresa Batista Neto, Monica Pozzato, Marie-Claude Blatter, Alex Ignatchenko, Nicole Redaschi, and Alan Bridge. Rhea, the reaction knowledgebase in 2022. *Nucleic Acids Research*, 50(D1):D693–D700, 2022. doi: 10.1093/nar/gkab1016.

Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P. Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E. Smidt, and Boris Kozinsky. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature Communications*, 13:2453, 2022. doi: 10.1038/s41467-022-29939-5.

Martin Buttenschoen, Garrett M. Morris, and Charlotte M. Deane. PoseBusters: AI-based docking methods fail to generate physically valid poses or generalise to novel sequences. *Chemical Science*, 15:3130–3139, 2024. doi: 10.1039/D3SC04185A.

Duanhua Cao, Geng Chen, Jiaxin Jiang, Jie Yu, Runze Zhang, Mingan Chen, Wei Zhang, Lifan Chen, Feisheng Zhong, Yingying Zhang, Chenghao Lu, Xutong Li, Xiaomin Luo, Sulin Zhang, and Mingyue Zheng. Generic protein–ligand interaction scoring by integrating physical prior knowledge and data augmentation modelling. *Nature Machine Intelligence*, 6:688–700, 2024. doi: 10.1038/s42256-024-00849-z.

Antje Chang, Lisa Jeske, Sandra Ulbrich, Julia Hofmann, Julia Koblitz, Ida Schomburg, Meina Neumann-Schaal, Dieter Jahn, and Dietmar Schomburg. BRENDA, the ELIXIR core data resource in 2021: new developments and updates. *Nucleic Acids Research*, 49(D1):D498–D508, 2021. doi: 10.1093/nar/gkaa1025.

Shuan Chen, Sunggi An, Ramil Babazade, and Yousung Jung. Precise atom-to-atom mapping for organic reactions via human-in-the-loop machine learning. *Nature Communications*, 15:2250, 2024. doi: 10.1038/s41467-024-46364-y.

Sunghwan Choi. Prediction of transition state structures of gas-phase chemical reactions via machine learning. *Nature Communications*, 14:1168, 2023. doi: 10.1038/s41467-023-36823-3.

Connor W. Coley, Wengong Jin, Luke Rogers, Timothy F. Jamison, Tommi S. Jaakkola, William H. Green, Regina Barzilay, and Klavs F. Jensen. A graph-convolutional neural network model for the prediction of chemical reactivity. *Chemical Science*, 10(2):370–377, 2019. doi: 10.1039/C8SC04228D.

Gabriele Corso, Hannes Stärk, Bowen Jing, Regina Barzilay, and Tommi S. Jaakkola. DiffDock: Diffusion steps, twists, and turns for molecular docking. In *International Conference on Learning Representations*, 2023.

Haiyang Cui, Yufeng Su, Tanner J. Dean, Tianhao Yu, Zhengyi Zhang, Jian Peng, Diwakar Shukla, and Huimin Zhao. Enzyme specificity prediction using cross-attention graph neural networks. *Nature*, 647:639–647, 2025. doi: 10.1038/s41586-025-09697-2.

Samir Darouich, Jacob W. Toney, Weiliang Luo, Johannes Kästner, Mathias Niepert, and Heather J. Kulik. Robust generative transition-state models for unseen chemistry. *Nature Computational Science*, 2026. doi: 10.1038/s43588-026-01034-5.

Chenru Duan, Yuanqi Du, Haojun Jia, and Heather J. Kulik. Accurate transition state generation with an object-aware equivariant elementary reaction diffusion model. *Nature Computational Science*, 3(12):1045–1055, 2023. doi: 10.1038/s43588-023-00563-7.

Chenru Duan, Guan-Horng Liu, Yuanqi Du, Tianrong Chen, Qiyuan Zhao, Haojun Jia, Carla P. Gomes, Evangelos A. Theodorou, and Heather J. Kulik. Optimal transport for generating transition states in chemical reactions. *Nature Machine Intelligence*, 7:615–626, 2025. doi: 10.1038/s42256-025-01010-0.

Kuzman Ganchev, João Graça, Jennifer Gillenwater, and Ben Taskar. Posterior regularization for structured latent variable models. *Journal of Machine Learning Research*, 11(67):2001–2049, 2010.

Mireia Garcia-Viloca, Jiali Gao, Martin Karplus, and Donald G. Truhlar. How enzymes work: analysis by modern rate theory and computer simulations. *Science*, 303(5655):186–195, 2004. doi: 10.1126/science.1088172.

Vladimir Gligorijević, P. Douglas Renfrew, Tomasz Kosciolek, Julia Koehler Leman, Daniel Berenberg, Tommi Vatanen, Chris Chandler, Bryn C. Taylor, Ian M. Fisk, Hera Vlamakis, Ramnik J. Xavier, Rob Knight, Kyunghyun Cho, and Richard Bonneau. Structure-based protein function prediction using graph convolutional networks. *Nature Communications*, 12:3168, 2021. doi: 10.1038/s41467-021-23303-9.

Colin A. Grambow, Lagnajit Pattanaik, and William H. Green. Reactants, products, and transition states of elementary chemical reactions based on quantum chemistry. *Scientific Data*, 7:137, 2020. doi: 10.1038/s41597-020-0460-4.

Akihide Hayashi, So Takamoto, Ju Li, Yuta Tsuboi, and Daisuke Okanohara. Generative model for constructing reaction path from initial to final states. *Journal of Chemical Theory and Computation*, 21(3):1292–1305, 2025. doi: 10.1021/acs.jctc.4c01397.

Esther Heid, Daniel Probst, William H. Green, and Georg K. H. Madsen. EnzymeMap: curation, validation and data-driven prediction of enzymatic reactions. *Chemical Science*, 14:14229–14242, 2023. doi: 10.1039/D3SC02048G.

Graeme Henkelman and Hannes Jónsson. Improved tangent estimate in the nudged elastic band method for finding minimum energy paths and saddle points. *The Journal of Chemical Physics*, 113(22):9978–9985, 2000. doi: 10.1063/1.1323224.

Graeme Henkelman, Blas P. Uberuaga, and Hannes Jónsson. A climbing image nudged elastic band method for finding saddle points and minimum energy paths. *The Journal of Chemical Physics*, 113(22):9901–9904, 2000. doi: 10.1063/1.1329672.

Wojciech Jaworski, Sara Szymkuć, Barbara Mikulak-Klucznik, Krzysztof Piecuch, Tomasz Klucznik, Michał Kaźmierowski, Jan Rydzewski, Anna Gambin, and Bartosz A. Grzybowski. Automatic mapping of atoms across both simple and complex chemical reactions. *Nature Communications*, 10:1434, 2019. doi: 10.1038/s41467-019-09440-2.

Wengong Jin, Connor W. Coley, Regina Barzilay, and Tommi Jaakkola. Predicting organic reaction outcomes with weisfeiler-lehman network. In *Advances in Neural Information Processing Systems*, volume 30, pp. 2607–2616, 2017.

Kjell Jorner, Anna Tomberg, Christoph Bauer, Christian Sköld, and Per-Ola Norrby. Organic reactivity from mechanism to machine learning. *Nature Reviews Chemistry*, 5:240–255, 2021. doi: 10.1038/s41570-021-00260-x.

Joonyoung F. Joung, Mun Hong Fong, Nicholas Casetti, Jordan P. Liles, Ne S. Dassanayake, and Connor W. Coley. Electron flow matching for generative reaction mechanism prediction. *Nature*, 645(8079):115–123, 2025. doi: 10.1038/s41586-025-09426-9.

Seonghwan Kim, Jeheon Woo, and Woo Youn Kim. Diffusion-based generative AI for exploring transition states from 2D molecular graphs. *Nature Communications*, 15:341, 2024. doi: 10.1038/s41467-023-44629-6.

Alexander Kroll, Sahasra Ranjan, Martin K. M. Engqvist, and Martin J. Lercher. A general model to predict small molecule substrates of enzymes based on machine and deep learning. *Nature Communications*, 14:2787, 2023. doi: 10.1038/s41467-023-38347-2.

Wei Lu, Qifeng Wu, Jixian Zhang, Jiahua Rao, Chengtao Li, and Shuangjia Zheng. TANKBind: Trigonometry-aware neural networks for drug-protein binding structure prediction. In *Advances in Neural Information Processing Systems*, volume 35, pp. 7236–7249, 2022. doi: 10.52202/068431-0525.

Wei Lu, Jixian Zhang, Weifeng Huang, Ziqiao Zhang, Xiangyu Jia, Zhenyu Wang, Leilei Shi, Chengtao Li, Peter G. Wolynes, and Shuangjia Zheng. DynamicBind: predicting ligand-specific protein–ligand complex structure with a deep equivariant generative model. *Nature Communications*, 15:1071, 2024. doi: 10.1038/s41467-024-45461-2.

Yufei Luo, Xiang Gu, and Jian Sun. Generative flow model on distance geometry for predicting transition states of chemical reactions. *Nature Communications*, 17:7497, 2026. doi: 10.1038/s41467-026-74101-0.

Marco Malatesta, Emanuele Fornasier, Martino Luigi Di Salvo, Angela Tramonti, Erika Zangelmi, Alessio Peracchi, Andrea Secchi, Eugenia Polverini, Gabriele Giachin, Roberto Battistutta, Roberto Contestabile, and Riccardo Percudani. One substrate many enzymes virtual screening uncovers missing genes of carnitine biosynthesis in human and mouse. *Nature Communications*, 15:3199, 2024. doi: 10.1038/s41467-024-47466-3.

Peter Mikhael, Itamar Chinn, and Regina Barzilay. CLIPZyme: Reaction-conditioned virtual screening of enzymes. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 35647–35663, 2024.

Xingang Peng, Shitong Luo, Jiaqi Guan, Qi Xie, Jian Peng, and Jianzhu Ma. Pocket2Mol: Efficient molecular sampling based on 3D protein pockets. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 17644–17655, 2022.

Jan Peters, Katharina Mülling, and Yasemin Altun. Relative entropy policy search. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*, pp. 1607–1612, 2010. doi: 10.1609/aaai.v24i1.7727.

Craig T. Porter, Gail J. Bartlett, and Janet M. Thornton. The catalytic site atlas: a resource of catalytic sites and residues identified in enzymes using structural data. *Nucleic Acids Research*, 32(suppl$_1$) : $D129 - -D133, 2004. doi : .$

Zhuoran Qiao, Weili Nie, Arash Vahdat, Thomas F. Miller, III, and Animashree Anandkumar. State-specific protein–ligand complex structure prediction with a multiscale deep generative model. *Nature Machine Intelligence*, 6:195–208, 2024. 10.1038/s42256-024-00792-z.

António J. M. Ribeiro, Gemma L. Holliday, Nicholas Furnham, Jonathan D. Tyzack, Katherine Ferris, and Janet M. Thornton. Mechanism and catalytic site atlas (M-CSA): a database of enzyme reaction mechanisms and active sites. *Nucleic Acids Research*, 46(D1):D618–D623, 2018. 10.1093/nar/gkx1012.

Reuven Y. Rubinstein and Dirk P. Kroese. *The Cross-Entropy Method: A Unified Approach to Combinatorial Optimization, Monte-Carlo Simulation and Machine Learning*. Springer, 2004. 10.1007/978-1-4757-4321-0.

Vern L. Schramm. Enzymatic transition states, transition-state analogs, dynamics, thermodynamics, and lifetimes. *Annual Review of Biochemistry*, 80:703–732, 2011. 10.1146/annurev-biochem-061809-100742.

Mathias Schreiner, Arghya Bhowmik, Tejs Vegge, Jonas Busk, and Ole Winther. Transition1x—a dataset for building generalizable reactive machine learning potentials. *Scientific Data*, 9:779, 2022. 10.1038/s41597-022-01870-w.

John Schulman, Sergey Levine, Pieter Abbeel, Michael I. Jordan, and Philipp Moritz. Trust region policy optimization. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 1889–1897, 2015.

Kristof T. Schütt, Pieter-Jan Kindermans, Huziel Enoc Sauceda Felix, Stefan Chmiela, Alexandre Tkatchenko, and Klaus-Robert Müller. SchNet: A continuous-filter convolutional neural network for modeling quantum interactions. In *Advances in Neural Information Processing Systems*, volume 30, pp. 991–1001, 2017.

Philippe Schwaller, Teodoro Laino, Théophile Gaudin, Peter Bolgar, Christopher A. Hunter, Costas Bekas, and Alpha A. Lee. Molecular transformer: A model for uncertainty-calibrated chemical reaction prediction. *ACS Central Science*, 5(9):1572–1583, 2019. 10.1021/acscentsci.9b00576.

Philippe Schwaller, Benjamin Hoover, Jean-Louis Reymond, Hendrik Strobelt, and Teodoro Laino. Extraction of organic chemistry grammar from unsupervised learning of chemical reactions. *Science Advances*, 7(15):eabe4166, 2021a. 10.1126/sciadv.abe4166.

Philippe Schwaller, Daniel Probst, Alain C. Vaucher, Vishnu H. Nair, David Kreutter, Teodoro Laino, and Jean-Louis Reymond. Mapping the space of chemical reactions using attention-based neural networks. *Nature Machine Intelligence*, 3:144–152, 2021b. 10.1038/s42256-020-00284-w.

Hans Martin Senn and Walter Thiel. QM/MM methods for biomolecular systems. *Angewandte Chemie International Edition*, 48(7):1198–1229, 2009. 10.1002/anie.200802019.

Yidong Song, Qianmu Yuan, Sheng Chen, Yuansong Zeng, Huiying Zhao, and Yuedong Yang. Accurately predicting enzyme functions through geometric graph learning on ESMFold-predicted structures. *Nature Communications*, 15:8180, 2024. 10.1038/s41467-024-52533-w.

Hannes Stärk, Octavian Ganea, Lagnajit Pattanaik, Regina Barzilay, and Tommi Jaakkola. EquiBind: Geometric deep learning for drug binding structure prediction. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 20503–20521, 2022.

Ping Tuo, Jiale Chen, and Ju Li. Flow matching for reaction pathway generation. *Nature Communications*, 17:8769, 2026. 10.1038/s41467-026-75654-w.

Marc W. van der Kamp and Adrian J. Mulholland. Combined quantum mechanics/molecular mechanics (QM/MM) methods in computational enzymology. *Biochemistry*, 52(16):2708–2728, 2013. 10.1021/bi400215w.

Xiaorui Wang, Xiaodan Yin, Dejun Jiang, Huifeng Zhao, Zhenxing Wu, Odin Zhang, Jike Wang, Yuquan Li, Yafeng Deng, Huanxiang Liu, Pei Luo, Yuqiang Han, Tingjun Hou, Xiaojun Yao, and Chang-Yu Hsieh. Multi-modal deep learning enables efficient and accurate annotation of enzymatic active sites. *Nature Communications*, 15:7348, 2024. 10.1038/s41467-024-51511-6.

Arieh Warshel, Pankaz K. Sharma, Mitsunori Kato, Yun Xiang, Hanbin Liu, and Mats H. M. Olsson. Electrostatic basis for enzyme catalysis. *Chemical Reviews*, 106(8):3210–3235, 2006. 10.1021/cr0503106.

Tianhao Yu, Haiyang Cui, Jianan Canal Li, Yunan Luo, Guangde Jiang, and Huimin Zhao. Enzyme function prediction using contrastive learning. *Science*, 379(6639):1358–1363, 2023. 10.1126/science.adf2465.

Eric C.-Y. Yuan, Anup Kumar, Xingyi Guan, Eric D. Hermes, Andrew S. Rosen, Judit Zádor, Teresa Head-Gordon, and Samuel M. Blau. Analytical ab initio hessian from a deep learning potential for transition state optimization. *Nature Communications*, 15:8865, 2024. 10.1038/s41467-024-52481-5.

Chengxin Zhang, Xi Zhang, Lydia Freddolino, and Yang Zhang. BioLiP2: an updated structure database for biologically relevant ligand–protein interactions. *Nucleic Acids Research*, 52(D1): D404–D412, 2024. 10.1093/nar/gkad630.

Fan Zhang, Tao Zeng, and Ruibo Wu. QM/MM modeling aided enzyme engineering in natural products biosynthesis. *Journal of Chemical Information and Modeling*, 63(16):5018–5034, 2023a. 10.1021/acs.jcim.3c00779.

Xujun Zhang, Odin Zhang, Chao Shen, Wanglin Qu, Shicheng Chen, Hanqun Cao, Yu Kang, Zhe Wang, Ercheng Wang, Jintu Zhang, Yafeng Deng, Furui Liu, Tianyue Wang, Hongyan Du, Langcheng Wang, Peichen Pan, Guangyong Chen, Chang-Yu Hsieh, and Tingjun Hou. Efficient and accurate large library ligand docking with KarmaDock. *Nature Computational Science*, 3(9):789–804, 2023b. 10.1038/s43588-023-00511-5.

A DATA UNITS, SUPERVISION, AND PROVENANCE

Experimental endpoint corpus. EnzymeMap supplies mapped reactants/products and net edits; BioLiP2/RCSB supplies holo evidence and coordinates. Exact enzyme and reactive-ligand linkage prevents EC-only matching or spectator ligands from being used as endpoint evidence. Directed reaction descriptions are deduplicated and forward/reverse views are merged into physical systems. The resulting broad split contains 3,021 training, 159 validation, 429 enzyme-OOD, 314 substrate-OOD, and 313 double-OOD physical systems. Accepted coordinate views are nested within those systems, not additional independent biological samples. The 1,992 training views belong to 1,771 training physical systems. Only present, graph-consistent coordinates contribute endpoint supervision.

Paired reaction cohort. The formal 171-system benchmark is a frozen paired-coordinate cache assembled for local refinement from public reaction records and the study's coordinate-generation/alignment pipeline; 112 additional public candidates lack at least one complete canonical endpoint record. It is distinct from the broad experimental-holo endpoint benchmark and is not presented as 171 pairs of independently crystallized ES/EP complexes. Each fold fits its own critic from the corresponding training records. No generic transition-region proposal is relabeled as a ground-truth enzyme transition state.

Enzyme groups use 40%-identity, 80%-coverage MMseqs2 clustering. Reaction groups are keyed by EC level 2 and the multiset of net edit types/bond-order changes. Connected components of the enzyme–reaction bipartite graph are assigned as whole units to five outer folds. This excludes overlap under those group definitions; it does not audit every upstream generator's pretraining data or prove a global maximum sequence similarity bound.

Table 4: Out-of-fold critic validation for the paired-reaction cohort. Values are recorded diagnostics, not the 863-view endpoint OOD results.

Fold	Train	Validation	Geometry acc.	Local acc.	AUPRC
0	106	30	93.3%	68.3%	72.6%
1	111	26	90.4%	80.8%	66.3%
2	115	22	93.2%	79.5%	68.6%
3	106	31	90.3%	71.0%	71.2%
4	106	31	85.5%	66.1%	62.2%

Training and sampling repetitions. All five paired-cohort critics use actor-optional training; supported roles are masked and the validation role gate requires sufficient support. Three ensemble members are jointly trained within each fold. The search seeds are 20260911, 20260912, and 20260913; they vary search randomness, not the fitted outer-fold ensemble. Three training seeds in the separate endpoint ablation table are a different repetition axis. The selected-checkpoint endpoint record gives AUPRC 83.04/88.03/80.25%; it must not be substituted for the three-seed aggregate 83.12/86.85/80.02%.

B REFERENCE POSTERIOR AND IMPLEMENTED INFORMATION BUDGETS

Let $\rho = p_{0,b}(\mathcal{F}_b) > 0$, and let $p_F = p_{0,b}\mathbf{1}_{\mathcal{F}_b}/\rho$. For any distribution supported on $\mathcal{F}_b$,

$$D_{\mathrm{KL}}(q\|p_{0,b}) = D_{\mathrm{KL}}(q\|p_F) - \log\rho. \tag{10}$$

Consequently, the alternative globally constrained problem $D_{\mathrm{KL}}(q\|p_{0,b}) \le \epsilon$ has nonempty feasible distributions only if $\epsilon \ge -\log\rho$. A hard support restriction can consume information budget before any utility improvement occurs. Eq. 2 uses a penalized reference objective and explicitly requires nonzero feasible prior mass.

If $Z_\eta = \int_{\mathcal{F}_b} p_{0,b}(x)e^{U_E(x)/\eta}\,dx$ is finite, then

$$\mathbb{E}_q[U_E] - \eta D_{\mathrm{KL}}(q\|p_{0,b}) = \eta\log Z_\eta - \eta D_{\mathrm{KL}}(q\|q_\eta^\star). \tag{11}$$

This proves the reference Gibbs solution without attributing a new optimizer theorem to CPT. It also clarifies that a delta-selected output is not the continuous posterior itself.

The practical Gaussian identity applies to consecutive *untruncated proposals* with identical covariance. The discrete weighting constraint in Eq. 5 is instead relative to the uniform measure on the current feasible particle indices. These are two different KL quantities. They are not a certificate of a fixed global KL between the final hard-filtered empirical distribution and the initial continuous Gaussian. Exact importance correction to a fixed $p_{0,b}$ is not supplied by uncorrected current-population weighting.

For $U' = aU + c$ with $a > 0$, the same normalized exponential weights follow from $\eta' = a\eta$ whenever the temperature solution is within its numerical bounds. This is the score-scale property of REPS. It neither proves global convergence nor implies superiority over CEM, which itself preserves elite ordering under positive affine score changes.

C SCORING, SUPPORT, AND OUTPUT DETAILS

The actor residual in Eq. 4 compares the nearest reacting-atom–actor distance of a candidate with the midpoint of the corresponding endpoint nearest distances. Distances are scaled by 1 Å and penalized with smooth-L1, with the configured transition scale 0.10 Å. The multiplier is 0.020. Missing actor evidence contributes zero. This residual is separate from the actor flags and role features input to the learned critic; removing one does not necessarily remove the other.

Same-system displacement negatives preserve internal ligand geometry, pocket chemistry, and annotation identity. Local and far displacement magnitudes are 0.5 and 2.5 Å for the default critic. Additional ranking terms compare native to chemistry permutation, supported actor reassignment, and supported role permutation, together with local-over-far ordering. The frozen wrong-pocket diagnostic uses size-matched non-native pockets with reaction-centre alignment; computing such a score in the forward pass does not mean its ranking loss contributes to optimization.

The severe-clash and edited-distance checks use fixed operational thresholds. The edited-pair interval is obtained from endpoint distances with the configured margin; edit progress and available actor contact may regress by at most 0.01 relative to the generic reference. These surrogate checks do not certify electronic structure, proton transfer, conservation along a complete mechanism, or catalytically productive binding.

The main run uses 64 antithetic particles per round under a common fixed refinement horizon. After refinement, a final proposal pool is sampled from the updated distribution before output filtering and selection. When fewer than eight particles are eligible, the preflight support-restoration phase applies bounded mean shifts up to its frozen restoration limit. Temperature bisection uses 64 iterations over $[10^{-6}, 10^{4}]$ and a 0.50-nat feasible-weight budget. The Gaussian step cap is 0.25 nats per reacting atom, with total mean RMS capped at $0.80b$. Operational convergence is evaluated from the round-wise traces using Eq. 7: proposal-mean RMS change at most 10^{-4} Å and change in weighted empirical utility at most 10^{-5} for two consecutive active rounds. The implementation records the per-round mean-step RMS, empirical-utility change, criterion indicator, and first convergence round. This is a diagnostic rather than solver-dependent early stopping in the matched experiments, so all CPT variants retain the same proposal budget. Insufficient support or flat utility disables a mean update and is not counted as convergence. Every returned candidate is rechecked and must meet the raw-prior utility floor; diverse hypotheses are selected at a minimum reacting-RMS separation of $0.10b$.

D BOND-EDIT CONSTRAINTS FROM MAPPED REACTION GRAPHS

For each mapped reactant/product pair, the reaction graph supplies bond additions, bond deletions, and bond-order changes. The union of atoms participating in these edits defines the reacting set. CPT uses the edited pairs to construct endpoint-derived distance intervals and to measure bond-edit progress. A proposal is rejected when its edited-pair geometry leaves the declared interval or when edit progress regresses beyond the frozen tolerance. These checks are applied together with the displacement and clash gates defined in the main text.

We use *bond-edit constraints* for this information throughout the paper. They encode the observed net graph transformation, not a complete sequence of elementary proton-transfer or intermediate states, and they are not a transition-state label. Catalytic actor/role annotations, when available, enter

the enzyme-context score and actor non-regression check; they are conceptually separate from the mapped bond edits.

E INDEPENDENT PHYSICAL CONTROLS AND CORRELATED PROFILES

The complete three-prior fixed-budget physical export uses CPT-REPS under the identifier `cpt`; this is why the cross-prior physical table is reported for REPS even though CEM is the stronger practical solver on learned utility. Raw-prior names, CEM, direct gradient, and context perturbations identify different references. The primary interval is attached to the paired system-level statistic for the named comparison; medians cannot be added or subtracted to reconstruct another comparison. In particular, REPS-minus-CEM and REPS-minus-raw summaries do not determine a CEM-minus-raw median or confidence interval.

A separate context-physical diagnostic compares native-context REPS with a matched wrong-pocket search on 104 usable React-OT systems. Native context has lower reaction-coordinate MAE on 49.0% of systems; the median native-minus-wrong-pocket difference is $+0.0007$ Å with 95% interval $[-0.0088, +0.0147]$. This scalar xTB proxy therefore does not establish an independent physical advantage for native context, and is kept separate from the endpoint context-discrimination results.

On React-OT, REPS and CEM are indistinguishable on the path-normal metric (median REPS-minus-CEM difference $+0.00042$ Å, 95% interval $[-0.00028, +0.00131]$), while CEM has lower reaction-coordinate MAE in Table 3. The path-normal control additionally shows why projected gradient cannot be judged from edited-pair error alone:

Table 5: React-OT path-normal control at 0.25 Å. Negative REPS-minus-reference RMS difference favors REPS. The archived metric is `ligand_heavy_path_normal_rmsd_A`.

Reference	n	Median difference (Å)	95% CI
Raw prior	110	$+0.00185$	$[-0.00066, +0.00400]$
CPT-CEM	110	$+0.00042$	$[-0.00028, +0.00131]$
Uniform	109	$+0.00103$	$[-0.00055, +0.00305]$
Projected gradient	109	-0.00408	$[-0.00615, -0.00056]$

The higher-coverage concerted scan requested 342 paths (two protonation variants for 171 systems); 318 were usable, 196 had an interior maximum, and 194 profiles could be paired. The 137/194 improvement fraction is a profile-level descriptive result. Among the 77 systems evaluable under both variants, 49 improve in both. Future inferential statements should retain the system as the bootstrap cluster rather than treating both variants as independent observations.

The historical heavy-atom error of 0.8495/0.8659 Å is computed as the square root of the mean squared Cartesian-coordinate differences in a fixed frame, not an optimally aligned atom-norm RMSD. Its lack of improvement is preserved without renaming it as a standard aligned pose metric.

F SECONDARY COVERAGE AND APPLICABILITY EVIDENCE

At the matched 172-profile downstream initializer stage, raw React-OT, CPT-REPS, and the xTB-path initializer obtain 67, 73, and 77 successful profiles, respectively. Their three-source union obtains 92/172. This describes complementary candidate coverage with up to three initializers, not a budget-matched improvement of one learned initializer. On the 155-profile common CEM subset, REPS/CEM/raw/path counts are 66/58/59/70; the three-source union is 83 and adding CEM raises it to 85. These summaries do not establish a same-budget top-k ranking advantage.

The observed prior applicability also limits transfer. In the 171-system cohort, heavy-atom size bins contain 20/75/67/9 reactions for ≤ 7, 8–14, 15–24, and ≥ 25 heavy atoms. REPS material-gain fractions decrease to 33.3% in the largest bin for both TS-DFM and MolGEN. OA-ReactDiff is evaluated only on its 29-system admitted coordinate-scale domain, not treated as a failed generator on all 171 systems.

A stratified semiempirical-to-DFT development screen does not furnish transition-state validation for this paper. Finite energies or a local negative-curvature direction alone do not establish a first-order saddle with the expected mode and two-sided connectivity. No activation-energy, kinetic, or experimentally validated enantioselectivity claim is made from these screening results.

G SEPARATELY EVALUATED ENDPOINT CONTEXT CORRECTION

The main endpoint context result uses a same-pocket chemistry permutation: pocket coordinates are fixed while residue/atom chemistry is permuted. This is distinct from replacing the whole pocket. The frozen endpoint critic did not directly optimize the size-matched foreign-pocket contrast. A later 1,500-step endpoint experiment adds this contrast but trades stronger context discrimination for reduced enzyme-OOD geometry AUPRC. The evaluated dual-critic variant standardizes each channel using validation scores and selects a context weight on a 0.05-spaced grid, subject to validation geometry and local-accuracy tolerances. The selected weight is 0.40. This variant has a separate model and selection protocol; its evaluation does not retroactively update the original out-of-fold search critics.

Table 6: Separate endpoint context follow-up on the same 863 OOD views. Single-checkpoint results, not means over three training seeds. "Base" denotes the geometry checkpoint used by this follow-up.

Metric	Enzyme	Substrate	Double
Views	287	274	302
Base geometry AUPRC (%)	83.04	88.03	80.25
Dual geometry AUPRC (%)	83.82	91.03	82.29
Base native-over-pocket (%)	49.48	57.66	56.29
Dual native-over-pocket (%)	51.92	64.96	58.94

The native-minus-foreign-pocket gap increases by 0.0598, 0.1135, and 0.0807 across the three splits, with reported intervals [0.0330,0.0869], [0.0826,0.1438], and [0.0499,0.1120]. These effects measure endpoint context discrimination. They are neither TS-geometry accuracy nor evidence that the dual critic has already improved all frozen search or physical results.