ZYMPRO: TRANSITION-AWARE GENERATIVE MODELING OF ENZYMATIC REACTION PROCESSES

Anonymous authors
Paper under double-blind review

ABSTRACT

Enzymatic reactions are fundamental to biological processes and underpin a wide range of biotechnological applications. During these reactions, transition states represent key molecular configurations that reflect enzyme-specific catalytic environments, which lower the activation barriers to product formation. However, despite extensive research on enzymatic reaction prediction, transition states remain underexplored in generative modeling, leaving unclear how they can support reaction prediction. Thus, we introduce ZymPro for generative modeling of enzymatic reaction processes through explicit transition-state generation. ZymPro employs energy-based process supervision from reaction paths and introduces a novel process-corrected diffusion that refines denoising predictions toward the transition state. The learned process knowledge is then transferred to product generation and enzyme compatibility ranking. Extensive experiments demonstrate that ZymPro achieves strong performance in enzyme compatibility ranking for unseen reaction rules and high accuracy in product generation for unseen products. The code is available at `https://anonymous.4open.science/r/ZymPro-26E4/`.

1 INTRODUCTION

Enzymes accelerate biochemical transformations by reshaping reaction pathways through their active-site environments. As illustrated in Fig. 1, reactants pass through a transition-state configuration before forming products. Unlike uncatalyzed reactions that follow intrinsic pathways, enzymes guide this transformation by providing specific active-site environments that stabilize the transition state. Through transition-state stabilization, enzymes reduce the activation free-energy barrier $\Delta G^{\ddagger}$, thereby accelerating the reaction without changing the overall reaction free-energy difference and equilibrium Warshel et al. (2006); Schramm (2011).

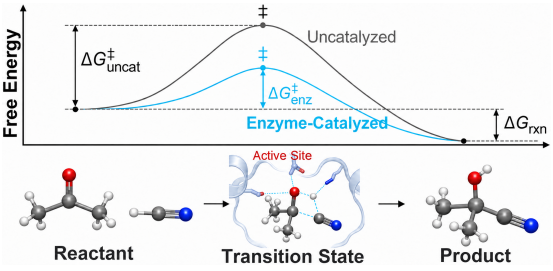

Figure 1: Enzyme catalysis through transition-state stabilization that lowers the activation free-energy.

The transition state represents the key molecular configuration in enzyme catalysis. Unlike uncatalyzed reactions that follow intrinsic reaction pathways, enzymes regulate biochemical transformations by selectively stabilizing transition-state configurations through enzyme-specific catalytic environments. Transition-state stabilization plays a critical role in catalytic efficiency, enzyme specificity, and product selectivity. For example, transition-state stabilization enhances catalytic efficiency in ketosteroid isomerase and Kemp eliminase and influences product selectivity in enzymatic Alder–ene reactions Fried et al. (2014); Blomberg et al. (2013); Ohashi et al. (2020). Therefore, transition states provide essential information for understanding enzymatic reaction processes.

Since transition states encode enzyme-dependent process information linking reactants to products, explicitly modeling them is valuable. Existing studies on enzymatic reaction modeling can be split into two areas. The first focuses on enzyme-reaction discrimination. CLIPZyme Mikhael et al. (2024) contrastively aligns enzymes with pseudo transition-state graphs, while ReactZyme

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.

Hua et al. (2024) models enzyme-reaction compatibility from multiple reaction views. However, these approaches focus on discrimination rather than modeling the transition-state mediated generative process. The second focuses on product generation. Sequence-to-sequence models such as Molecular Transformer Schwaller et al. (2019); Lee et al. (2019) directly map reactants to products, but overlook enzyme-mediated catalysis and the transition states underlying product formation. Initial efforts such as TSDiff and GoFlow Kim et al. (2024); Galustian et al. (2025) incorporate transition states into generative reaction modeling. However, their original formulations require product information and ignore enzyme-specific catalytic effects. Therefore, we aim to investigate transition-aware generative modeling of enzymatic reaction processes.

To capture the catalytic processes underlying enzymatic transformations, we propose to investigate transition-aware generative modeling of enzymatic reaction processes. This objective raises two major challenges: *First*, accurate transition-state annotations are essential for modeling the reaction process connecting reactants and products. However, transition states exist only fleetingly and cannot be isolated as stable species, making their experimental characterization difficult and reliable supervision scarce. Therefore, a key challenge is how to learn transition-aware generative processes in the absence of experimental transition-state data. *Second*, although transition states provide critical information connecting enzyme-specific catalytic environments with product formation, how to exploit such process-level knowledge in downstream tasks remains unclear.

To address these challenges, we propose ZymPro, which uses transition-state and reaction-path information as process-level supervision for enzyme-conditioned prediction. To model the reaction process connecting reactants and products, ZymPro introduces Process-Corrected Residual Diffusion, which jointly trains a transition-state predictor and a path-guided Corrector. Additionally, energy-based process supervision constructed from computational transition states and their associated reaction paths addresses the scarcity of experimental annotations (Peters et al., 2004; Bannwarth et al., 2019). To support downstream prediction, ZymPro transfers the learned reaction-aware representations and uses generated transition states to guide product generation and TS-supervised representations for enzyme compatibility ranking. In summary, our main contributions are as follows:

- We investigate transition-aware generative modeling of enzymatic reaction processes through explicit transition-state generation from the enzyme and reactants.
- We propose ZymPro, which combines pairwise enzyme–reactant interactions with energy-based process supervision to generate transition states with Corrector refinement, and transfers process knowledge in task-specific forms to product generation and enzyme compatibility ranking.
- Extensive experiments demonstrate the effectiveness of ZymPro in transition-state generation, product generation, and enzyme compatibility ranking, with ablations validating the benefits of process supervision and transition-state evidence.

2 PRELIMINARIES

This section introduces transition states, preliminary experiments, and enzymatic reaction modeling.

2.1 TRANSITION STATES IN ENZYME-CATALYZED REACTIONS

Enzyme-Catalyzed Reaction Processes. An enzyme $\mathcal{E}$ catalyzes the conversion of molecular input $\mathcal{R}$ (reactants and reagents) into products $\mathcal{P}$. In Fig. 1 (bottom), the enzyme first recognizes and binds substrates, which then adopt a reactive arrangement and pass through a transition state to form products. Within this enzyme-specific catalytic environment, the spatial arrangement and physicochemical properties of active- and binding-site residues constrain substrate positioning and molecular interactions, contributing to catalytic specificity (Cui et al., 2025).

Transition States in Enzyme-Catalyzed Reactions. A transition state $\mathcal{T}$ is a transient, high-energy configuration along the conversion of $\mathcal{R}$ to $\mathcal{P}$. Its geometry and electronic structure reflect the extent of bond breaking and formation and the associated charge redistribution (Schramm, 2011). In Fig. 1 (top), interactions in the catalytic environment of $\mathcal{E}$ preferentially stabilize $\mathcal{T}$ relative to $\mathcal{R}$, lowering the activation free-energy barrier $\Delta G^{\ddagger}$ and facilitating the reaction (Fried et al., 2014). Differential stabilization of competing transition states can favor particular pathways and products, contributing to catalytic selectivity (Ohashi et al., 2020). Transition-state information thus links enzyme-specific catalytic environments to the chemical changes underlying product formation.

2.2 Limitations of Generative Models in Enzymic Reaction Processes

Existing methods commonly formulate enzyme-conditioned product prediction as an endpoint mapping $f : (\mathcal{E}, \mathcal{R}) \mapsto \mathcal{P}$ (Wang et al., 2026; Kim et al., 2025). However, these models may primarily learn reactant-product associations without capturing the enzyme-mediated catalytic mechanism. We conduct the following experiments to demonstrate the limitations of representative existing methods in modeling enzyme-catalyzed reaction processes.

- **Sensitivity to Input Enzymes:** A model that overfits reactant-product associations may be insensitive to enzyme inputs. In contrast, a model that captures the enzyme's role in catalysis should produce predictions that reflect enzyme-dependent differences in reaction outcomes. We therefore replace the input enzyme and measure the fraction of unchanged product predictions.
- **Enzyme Discrimination:** A generative model that captures the enzyme's role in the reaction process should assign higher conditional probabilities $p(\mathcal{P} \mid \mathcal{E}, \mathcal{R})$ to enzymes compatible with the observed reactant-to-product transformation. Given the reactants and target product, we rank candidate enzymes by these probabilities and evaluate Top-1 accuracy, defined as the fraction of reactions for which an annotated enzyme ranks first.

We evaluate MetaReact, CYP-MAP, CyProduct, and EC3-MT (Wang et al., 2026; Kim et al., 2025; Tian et al., 2021; Schwaller et al., 2019) on reactions from the CYP-MAP benchmark (Kim et al., 2025). Detailed experimental settings are provided in Appendix A.1. As shown in Fig. 2, we make two observations: (**i**) Replacing the enzyme leaves the predicted product unchanged in at least 20% of cases across all four models, exceeding 60% for EC3-MT. (**ii**) Given the reactants and target product, all four models rank a matched enzyme first in fewer than 50% of cases, with EC3-MT close to the random guess. These observations indicate limitations of existing generative models in capturing enzyme-catalyzed reaction processes.

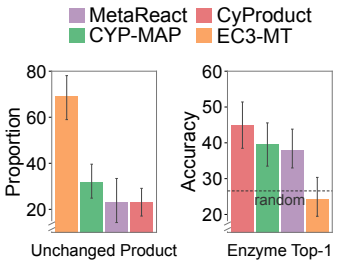

Figure 2: Enzyme-Swap Diagnostics.

2.3 Problem Definition

Notations. An enzyme with l residues is represented as $\mathcal{E} = (\mathbf{Q}, \mathbf{X}, \mathbf{a})$, where $\mathbf{Q} \in \mathbb{R}^{l \times d_E}$ denotes residue features, $\mathbf{X} \in \mathbb{R}^{l \times 3}$ denotes Cα coordinates, and $\mathbf{a} \in \{0, 1\}^l$ indicates active- and binding-site residues. A reactant-side input with n atoms is represented as $\mathcal{R} = (\mathbf{q}, \mathbf{x}, \mathbf{s}^R)$, where $\mathbf{q} \in \mathbb{R}^{n \times d_R}$ and $\mathbf{x} \in \mathbb{R}^{n \times 3}$ denote atom features and coordinates, respectively, and $\mathbf{s}^R = (s_1^R, \ldots, s_{L_R}^R)$ is the reactant-side SMILES token sequence containing the reactants and any recorded reagents. A transition state is denoted by $\mathcal{T}$ and represented by its atom-pair distance matrix $D^{TS} \in \mathbb{R}^{n \times n}$. Following prior works on reaction product generation (Schwaller et al., 2019; Tu & Coley, 2022; Sagawa & Kojima, 2025), we represent the product $\mathcal{P}$ by the SMILES tokens $\mathbf{s}^P = (s_1^P, \ldots, s_{L_P}^P)$. For example, the product aspirin is represented as `CC(=O)Oc1ccccc1C(=O)O`.

Problem 1. Given an enzyme $\mathcal{E}$ and reactants $\mathcal{R}$, transition-aware generative modeling of enzymatic reaction processes aims to learn the following conditional distribution

$$p_{\phi,\theta}(\mathcal{P}, \mathcal{T} \mid \mathcal{E}, \mathcal{R}) = p_\phi(\mathcal{T} \mid \mathcal{E}, \mathcal{R})\, p_\theta(\mathcal{P} \mid \mathcal{E}, \mathcal{R}, \mathcal{T}). \tag{1}$$

where p_ϕ models transition state generation conditioned on the enzyme and p_θ models product generation conditioned on the enzyme, reactants, and transition state.

3 Methodology

As illustrated in Fig. 3, ZymPro explicitly models enzyme-dependent transition states for downstream prediction. Two technical challenges need to be addressed: (i) how to learn transition state generation conditioned on the enzyme without experimental transition-state annotations, and (ii) how to exploit the learned process knowledge for product generation and enzyme compatibility ranking. To address these challenges, ZymPro introduces Process-Corrected Residual Diffusion for transition-state generation. Additionally, energy-based process supervision from computational transition states and reaction paths addresses the scarcity of experimental transition-state annotations. To support downstream prediction, the learned reaction-aware representations and generated transition states then guide product generation and enzyme compatibility ranking.

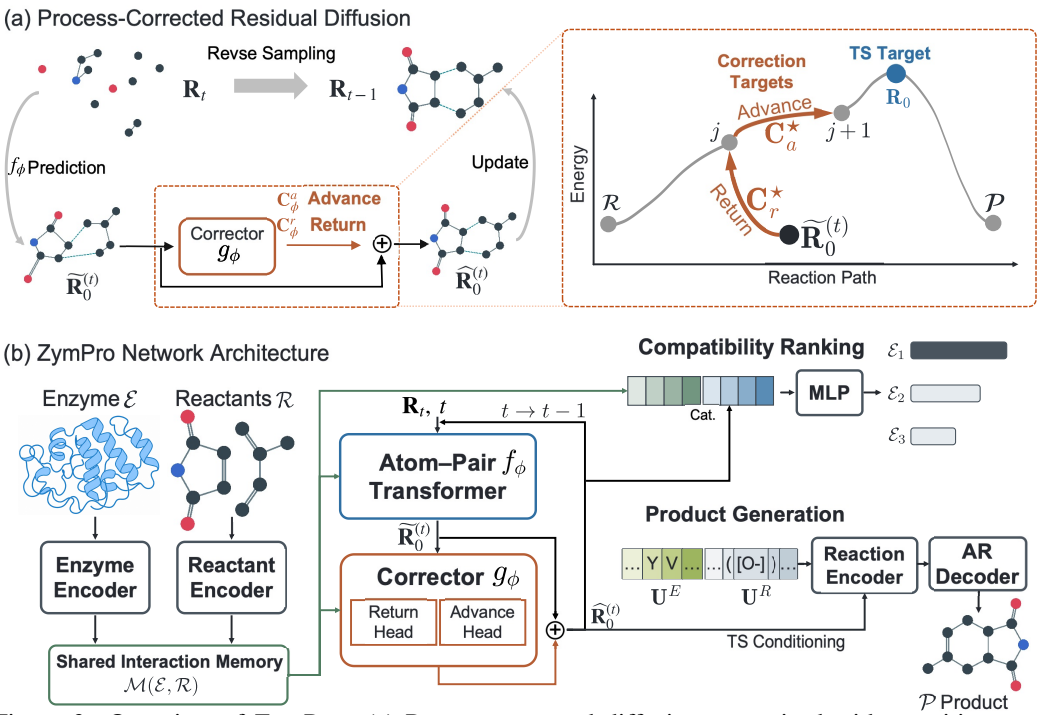

Figure 3: Overview of ZymPro. (a) Process-corrected diffusion supervised with transition-state and correction targets. Structures schematically depict reactant distances with added residuals. (b) Network architecture combining encoding, transition-state generation, and downstream tasks.

3.1 Transition-State Generation under Energy-Based Process Supervision

With limited experimental transition-state annotations, we introduce energy-based process supervision from computationally constructed reaction paths. This supervision guides the training and correction of the enzyme-conditioned diffusion model to learn enzyme-conditioned reaction geometry from computational paths. We next describe the construction of this supervision and the framework for transition-state generation.

Construction of Energy-Based Process Supervision. To address the scarcity of experimental transition-state annotations, we construct computational reaction paths connecting reactants and products. For each reaction i, FlowDock (Morehead & Cheng, 2025) predicts both endpoint poses in the enzyme pocket, which are aligned to a common atom order. Double-ended molecularGSM (Peters et al., 2004; Zimmerman, 2013) grows the path inward from these endpoints, initializing intermediate configurations $\mathbf{X}_i^{(k)}$. GFN2-xTB (Bannwarth et al., 2019) supplies energies and gradients to refine the reaction path, illustrated locally by

$$\mathbf{X}_i^{(k)} \leftarrow \mathbf{X}_i^{(k)} - \eta \, \mathrm{Proj}_{\perp \tau_{i,k}} \left[\nabla_{\mathbf{X}} E_{\mathrm{xTB}} \left(\mathbf{X}_i^{(k)} \right) \right], \tag{2}$$

where $E_{\mathrm{xTB}}(\mathbf{X})$ is the joint energy of the reacting molecules and active-site context pocket; η is a step size, and the projection retains gradients perpendicular to the path tangent $\tau_{i,k}$.

After optimization and validity filtering, we obtain the ordered reaction frames and select the highest-energy internal frame as the approximate transition-state target (Appendix A.2):

$$\Gamma_i = \left(\mathbf{X}_i^{(0)}, \mathbf{X}_i^{(1)}, \ldots, \mathbf{X}_i^{(K)} \right), \quad \mathbf{X}_i^{TS} = \underset{\mathbf{X} \in \{\mathbf{X}_i^{(1)}, \ldots, \mathbf{X}_i^{(K-1)}\}}{\arg\max} E_{\mathrm{xTB}}(\mathbf{X}), \tag{3}$$

where $\mathbf{X}_i^{(k)}$ contains the atomic coordinates of frame k, with reactant and product endpoints at $k = 0$ and $k = K$. The selected target supervises denoising, while intermediate frames supervise path correction.

Process-Corrected Residual Diffusion. We parameterize the transition-state generator $p_\phi(\mathcal{T} \mid \mathcal{E}, \mathcal{R})$ in Eq. 1 using residual diffusion conditioned on the enzyme and reactants. Since most atom-pair distances change little during a reaction, we model the clean residual $\mathbf{R}_0 = (\mathbf{D}^{TS} - \mathbf{D}^R)/s_\Delta$,

where $\mathbf{D}^R$ and $\mathbf{D}^{TS}$ are the reactant and transition-state distance matrices, and s_Δ is a fixed scale. The forward process adds Gaussian noise:

$$\mathbf{R}_t = \sqrt{\bar{\alpha}_t}\,\mathbf{R}_0 + \sqrt{1-\bar{\alpha}_t}\,\boldsymbol{\epsilon}, \qquad \boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I}), \qquad \bar{\alpha}_t = \prod_{u=1}^{t}(1-\beta_u), \tag{4}$$

where $\{\beta_t\}$ is a fixed variance schedule. Noise is applied to valid atom pairs while preserving symmetry and a zero diagonal.

For reverse sampling, we adopt an x_0-prediction parameterization (Salimans & Ho, 2022), directly estimating $\mathbf{R}_0$ from $\mathbf{R}_t$. However, these intermediate estimates may deviate from the reaction path despite receiving endpoint supervision. We therefore use the energy-based process frames Γ_i in Eq. (3) to supervise a Corrector that refines each clean prediction. It learns a return displacement toward the reaction path and an advance displacement along it toward the transition state (Fig. 3a):

$$\widetilde{\mathbf{R}}_0^{(t)} = f_\phi(\mathbf{R}_t, t; \mathcal{M}), \quad (\mathbf{C}_\phi^r, \mathbf{C}_\phi^a) = g_\phi(\widetilde{\mathbf{R}}_0^{(t)}, t; \mathcal{M}), \quad \widehat{\mathbf{R}}_0^{(t)} = \widetilde{\mathbf{R}}_0^{(t)} + \mathbf{C}_\phi^r + \mathbf{C}_\phi^a. \tag{5}$$

Here, $\mathcal{M} = \mathcal{M}(\mathcal{E}, \mathcal{R})$ encodes the enzyme and reactants, and $\mathbf{C}_\phi^r, \mathbf{C}_\phi^a$ are bounded corrections in residual-distance space. The corrected prediction determines the noise remaining in $\mathbf{R}_t$. Following deterministic DDIM sampling (Song et al., 2021), we combine this prediction with the inferred noise rescaled to the next timestep:

$$\mathbf{R}_{t-1} = \sqrt{\bar{\alpha}_{t-1}}\,\widehat{\mathbf{R}}_0^{(t)} + \sqrt{\frac{1-\bar{\alpha}_{t-1}}{1-\bar{\alpha}_t}}\left(\mathbf{R}_t - \sqrt{\bar{\alpha}_t}\,\widehat{\mathbf{R}}_0^{(t)}\right). \tag{6}$$

For brevity, $t-1$ denotes the next step in the reverse sampling sequence, with coefficients evaluated at the corresponding scheduled timesteps. Starting from Gaussian noise, repeated updates yield the final clean prediction $\widehat{\mathbf{R}}_0$, giving $\widehat{\mathbf{D}}^{TS} = \mathbf{D}^R + s_\Delta \widehat{\mathbf{R}}_0$.

Network Architecture. To incorporate the catalytic environment, prediction and correction share an enzyme–reactant interaction memory $\mathcal{M}$ (Fig. 3). It combines atom, residue, and pair features from Uni-Mol2 (Ji et al., 2024) and UniZyme (Li et al., 2025), which encode reactant geometry and enzyme sequence–structure information. Conditioned on $\mathcal{M}$ and timestep t, an atom–pair Transformer f_ϕ maps $\mathbf{R}_t$ to the clean residual $\widetilde{\mathbf{R}}_0^{(t)}$. The Corrector g_ϕ uses this estimate and the same conditions to predict return and advance corrections through two branches.

Learning Objective. We construct endpoint and correction targets from the reaction frames Γ_i in Eq. (3). We convert each frame $\mathbf{X}_i^{(k)} \in \Gamma_i$ into residual distances using $(\mathbf{D}(\mathbf{X}_i^{(k)}) - \mathbf{D}_i^R)/s_\Delta$, where $\mathbf{D}(\cdot)$ computes pairwise distances. The selected frame $\mathbf{X}_i^{TS}$ defines $\mathbf{R}_0$. We interpolate these residual frames in reaction order and find the position closest to $\widetilde{\mathbf{R}}_0^{(t)}$, denoted by $\mathbf{R}_\Gamma^{(j)}$. Together with the next position toward the transition state, $\mathbf{R}_\Gamma^{(j+1)}$, it defines the correction targets:

$$\mathbf{C}_r^\star = \mathbf{R}_\Gamma^{(j)} - \widetilde{\mathbf{R}}_0^{(t)}, \qquad \mathbf{C}_a^\star = \mathbf{R}_\Gamma^{(j+1)} - \mathbf{R}_\Gamma^{(j)}. \tag{7}$$

We jointly train the predictor and Corrector with

$$\mathcal{L}_{TS} = \mathbb{E}_{i,t,\boldsymbol{\epsilon}}\left[\left\|\widehat{\mathbf{R}}_0^{(t)} - \mathbf{R}_0\right\|_F^2 + \lambda_c\big(\ell(\mathbf{C}_\phi^r, \mathbf{C}_r^\star) + \ell(\mathbf{C}_\phi^a, \mathbf{C}_a^\star)\big) + \mathcal{L}_{\text{aux}}\right], \tag{8}$$

where ℓ is a robust regression loss and $\mathcal{L}_{\text{aux}}$ collects structural and energy-profile regularizers. Correction targets are detached during training, with zero advance at the transition state. Reference frames are not required at inference.

3.2 Product Generation Conditioned on Transition States

We model $p_\theta(\mathcal{P} \mid \mathcal{E}, \mathcal{R}, \mathcal{T})$ in Eq. 1 using a Transformer encoder–decoder. To guide product generation with reaction geometry, we incorporate the predicted TS distances and their changes from the reactants as pairwise attention biases. Formally, let $\widehat{\mathbf{D}}^{TS}$ represent the transition state $\widehat{\mathcal{T}}$ generated in Section 3.1. We construct the attention bias as

$$\mathbf{B}^{TS} = \mathrm{MLP}_{\text{abs}}\Big(\mathbf{W}\ \mathrm{RBF}(\widehat{\mathbf{D}}^{TS})\,\mathbf{W}^\top\Big) + \mathrm{MLP}_{\text{rel}}\Big(\mathbf{W}\ \mathrm{RBF}(\widehat{\mathbf{D}}^{TS} - \mathbf{D}^R)\,\mathbf{W}^\top\Big), \tag{9}$$

where RBF denotes Gaussian basis expansion and $\mathbf{W}$ is the normalized token–atom alignment matrix, applied to each feature channel to map atom-pair features to reactant token pairs. The two MLPs convert absolute-distance and distance-change features into additive attention biases.

Given reactant token embeddings $\mathbf{U}^R$ and selected enzyme residue representations $\mathbf{U}^E$, product generation is formulated as

$$\mathbf{H}_{\text{enc}} = f_\theta(\mathbf{U}^R, \mathbf{U}^E; \mathbf{B}^{TS}), \quad p_\theta(\mathcal{P} \mid \mathcal{E}, \mathcal{R}, \widehat{\mathcal{T}}) = \prod_{\ell=1}^{L_P} p_\theta(s_\ell^P \mid \mathbf{s}_{<\ell}^P, \mathbf{H}_{\text{enc}}). \tag{10}$$

Here, f_θ jointly encodes reactant tokens and catalytic-site and pocket residues, with $\mathbf{B}^{TS}$ added to self-attention logits between reactant tokens. The decoder cross-attends to $\mathbf{H}_{\text{enc}}$ to generate the product SMILES sequence, where s_ℓ^P is its ℓ-th token and L_P is its length. We train the model by minimizing the following objective function

$$\mathcal{L}_P = -\sum_{\ell=1}^{L_P} \log p_\theta(s_\ell^P \mid \mathbf{s}_{<\ell}^P, \mathbf{H}_{\text{enc}}). \tag{11}$$

The product encoder and decoder are pretrained on USPTO-MIT (Schwaller et al., 2019) and fine-tuned on EnzymeMap with enzyme and TS conditioning, while the interaction encoder and TS generator remain frozen. Appendix A.5 provides implementation details.

3.3 Enzyme Compatibility Ranking with Transition-State Evidence

Enzyme–reactant interactions characterize the initial complex, while predicted transition states provide complementary evidence about the reaction process. We therefore augment enzyme–reactant scoring $s(\mathcal{E}, \mathcal{R})$ with predicted transition-state evidence extracted by one denoising and correction step at the reactant anchor for ranking. We reuse the enzyme and reactant representations $\mathbf{H}_E, \mathbf{H}_R$ from Section 3.1. We obtain the predicted transition state $\widehat{\mathcal{T}}$ from the enzyme $\mathcal{E}$ and reactants $\mathcal{R}$ using the model in Eq. (5). An MLP-based scorer combines their pooled features to compute the compatibility score:

$$s(\mathcal{E}, \mathcal{R}, \widehat{\mathcal{T}}) = \sigma\Big(\text{MLP}(\mathbf{H}_E, \mathbf{H}_R, \widehat{\mathcal{T}})\Big), \tag{12}$$

where σ is sigmoid and MLP sums direct and transition-state scores. For fixed reactants, candidate enzymes are ranked by compatibility score. Appendix A.5 details feature extraction and training.

4 Experiments

In this section, we conduct experiments to answer the following research questions:

- **RQ1:** How well does ZymPro rank candidate enzymes under the rule-disjoint setting?
- **RQ2:** How accurately does ZymPro generate products under the product-disjoint setting?
- **RQ3:** Does process-corrected generation follow the reaction process captured by computational reference paths?

4.1 Experimental Setup

Dataset. We use EnzymeMap (Heid et al., 2023) as the source of enzymatic reactions and enzyme annotations. The processed records cover 491 valid reaction rules and form the basis of our enzyme compatibility ranking and product generation datasets. To provide supervision for transition-state generation, we additionally compute reaction paths and transition-state targets for eligible reactions, following Section 3.1. For compatibility ranking, we construct negative pairs by replacing the original enzyme through EC-stratified sampling (Appendix A.6). We also use USPTO-MIT (Schwaller et al., 2019) as a broader reaction-SMILES corpus to pretrain the product model before fine-tuning it on EnzymeMap.

Table 1: Task-specific dataset statistics.

	Rule-disjoint		Product-disjoint	
	TS Generation	Enzyme Ranking	Product Generation	Pretraining
Train	28,974	87,849	110,646	396,539
Val	3,141	6,464	10,313	29,084
Test	3,142	5,255	9,519	38,778

Evaluation. We evaluate generalization to unseen reaction rules and products under two settings. Table 1 reports the final training, validation, and test set sizes after task-specific filtering. Dataset construction and leakage controls are detailed in Appendix A.6.

- **Rule-Disjoint Setting:** We group reactions by rule for transition-state generation and enzyme compatibility ranking, so validation/test rules are absent from training. Shared records retain the same rule-disjoint assignment across transition-state generation and ranking. The training, validation, and test splits cover 178, 84, and 75 rules for compatibility ranking and 168, 100, and 108 rules for transition-state generation.
- **Product-Disjoint Setting:** We split reactions into training, validation, and test sets with no overlap in products, covering 57, 33, and 34 EC families (L2). Products are canonicalized while preserving stereochemistry and all product components. Reactions whose products appear in the EnzymeMap validation or test set are excluded from USPTO-MIT pretraining.
- **Evaluation Metrics:** For enzyme compatibility ranking, we report candidate ranking metrics (**Top-1**, **MRR**, **Hits@2**), pair-level discrimination metrics (**ROC-AUC**, **PR-AUC**), and early-retrieval metrics (**BEDROC** with $\alpha = 85, 20$ and **EF** at 5% and 10%) (Truchon & Bayly, 2007). Product generation is assessed using **Top-1**, **Top-3**, and **Top-5** exact-match accuracy.

Baseline Methods. We compare seven enzyme-reaction models covering embedding similarity, feature fusion, and cross-attention (Table 2), and seven product generators covering sequence-to-sequence, graph-to-sequence, and graph generation (Table 3). For compatibility ranking, baselines requiring product inputs are adapted to use only reactants and enzymes. Task-specific adaptations and additional evaluation of transition-state generation are detailed in Appendices A.7, A.8, and A.9.

Implementation Details. The transition-state generator uses 25-step DDIM sampling over a 1,000-step diffusion schedule and remains frozen downstream. Compatibility ranking extracts evidence using one denoising operation followed by one correction operation and trains a randomly initialized MLP. We select checkpoints on validation data; implementation details are in Appendices A.4 and A.5.

Table 2: Enzyme compatibility ranking on 5,255 groups of reactions

Model	Top-1	MRR	Hits@2	ROC-AUC	PR-AUC	B_{85}	B_{20}	$EF_{0.05}$	$EF_{0.1}$
CLIPZyme (Mikhael et al., 2024)	$39.49_{1.27}$	$63.74_{0.73}$	$69.59_{1.30}$	$65.37_{1.03}$	$37.05_{1.65}$	$51.54_{2.36}$	$43.42_{1.94}$	$1.75_{0.11}$	$1.65_{0.09}$
ReactZyme (Hua et al., 2024)	$47.92_{1.30}$	$68.06_{0.82}$	$70.82_{1.13}$	$65.81_{0.98}$	$40.59_{1.76}$	$60.01_{2.55}$	$52.03_{2.03}$	$2.30_{0.11}$	$1.94_{0.09}$
EnzRank (Upadhyay et al., 2023)	$36.60_{1.23}$	$61.43_{0.87}$	$66.38_{1.30}$	$62.05_{0.92}$	$34.08_{1.56}$	$46.21_{2.82}$	$41.02_{1.90}$	$1.74_{0.12}$	$1.59_{0.08}$
FusionESP (Du et al., 2025)	$\underline{54.00}_{1.34}$	$\underline{71.95}_{0.84}$	$\underline{75.76}_{1.18}$	$68.65_{1.02}$	$\underline{50.68}_{1.78}$	$\underline{81.96}_{1.92}$	$\underline{67.19}_{1.85}$	$\underline{2.97}_{0.12}$	$\underline{2.46}_{0.09}$
ESP (Kroll et al., 2023)	$44.35_{1.42}$	$67.19_{0.84}$	$73.49_{1.13}$	$63.86_{0.90}$	$34.47_{1.50}$	$47.29_{2.74}$	$39.31_{1.77}$	$1.56_{0.12}$	$1.54_{0.09}$
ProSmith (Kroll et al., 2024)	$51.66_{1.19}$	$71.03_{0.82}$	$75.45_{1.04}$	$\underline{72.25}_{0.94}$	$44.15_{1.59}$	$58.50_{2.62}$	$58.27_{1.79}$	$2.34_{0.12}$	$2.34_{0.09}$
EZSpecificity (Cui et al., 2025)	$51.56_{1.30}$	$70.73_{0.78}$	$74.72_{1.22}$	$70.59_{0.99}$	$46.53_{1.66}$	$56.99_{2.61}$	$59.13_{1.87}$	$2.56_{0.13}$	$2.35_{0.08}$
ZymPro	$65.11_{1.25}$	$79.80_{0.79}$	$85.16_{0.96}$	$81.58_{0.90}$	$62.88_{1.76}$	$86.60_{2.07}$	$76.91_{1.69}$	$3.25_{0.09}$	$3.09_{0.07}$

4.2 Enzyme Compatibility Ranking

To answer **RQ1**, we compare ZymPro with seven enzyme-reaction baselines under the rule-disjoint setting, using identical data splits. Table 2 reports overall ranking performance. Figure 4 compares ZymPro with the strongest baseline in each of 21 reaction rule classes containing at least 20 reactions. ZymPro outperforms the strongest baseline in 16 of the 21 classes. In Table 2, ZymPro performs best on all nine metrics, with Top-1 of 65.11% versus FusionESP's 54.00%. Its ROC-AUC is 81.58% versus ProSmith's 72.25%, and PR-AUC is 62.88% versus FusionESP's 50.68%, indicating better discrimination between enzyme-reactant pairs. Its $BEDROC_{85}$ of 86.60% and $EF_{0.05}$ of 3.25 further indicate that compatible pairs are concentrated near the top of the ranked list. Higher MRR and Hits@2 also indicate accurate enzyme selection from candidate groups under unseen reaction rules.

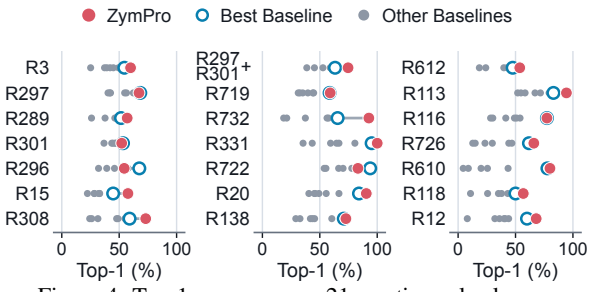

Figure 4: Top-1 scores across 21 reaction rule classes.

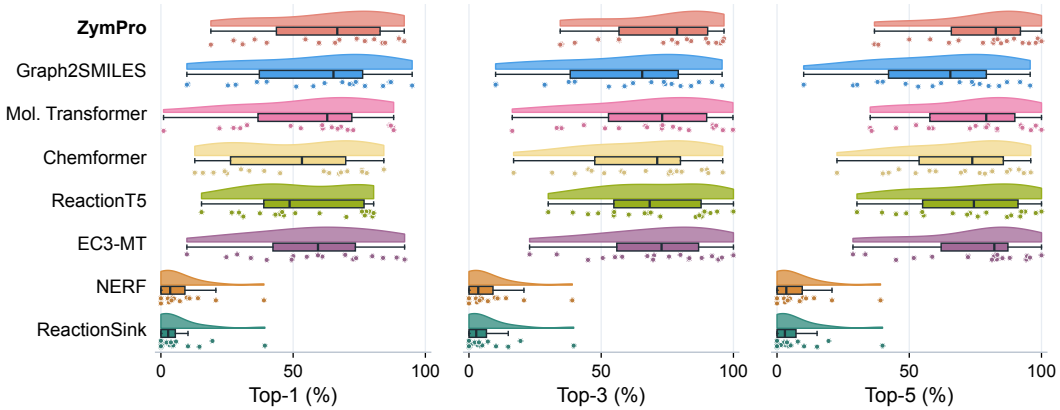

Figure 5: Exact-match accuracies for product generation across 18 EC classes.

4.3 Product Generation

To answer **RQ2**, we compare ZymPro with seven product generators on 9,519 product-disjoint test reactions. All models are pretrained from scratch on the same USPTO-MIT corpus and fine-tuned on the same data split; Table 3 summarizes the overall results, and Figure 5 shows the distributions of Top-1, Top-3, and Top-5 accuracies across 18 EC (L2) classes with at least 20 reactions. ZymPro achieves the best Top-1 accuracy at 53.17%, versus Graph2SMILES's 49.30%. It also exceeds EC3-MT in Top-3 (67.22% vs. 65.02%) and Top-5 (72.44% vs. 71.92%).

Table 3: Product generation on 9,519 reactions.

Model	Top-1 ↑	Top-3 ↑	Top-5 ↑
Graph2SMILES (Tu & Coley, 2022)	$\underline{49.30}_{1.05}$	$52.97_{1.05}$	$59.79_{1.06}$
Mol. Transformer (Schwaller et al., 2019)	$47.49_{1.05}$	$62.70_{1.03}$	$68.65_{0.99}$
Chemformer (Irwin et al., 2022)	$42.58_{1.02}$	$62.29_{1.02}$	$68.19_{1.00}$
ReactionT5 (Sagawa & Kojima, 2025)	$45.02_{0.92}$	$64.63_{0.88}$	$68.58_{0.85}$
EC3-MT (Schwaller et al., 2019)	$47.22_{0.99}$	$\underline{65.02}_{0.96}$	$\underline{71.92}_{0.92}$
NERF (Bi et al., 2021)	$10.33_{0.55}$	$10.35_{0.55}$	$10.40_{0.55}$
ReactionSink (Meng et al., 2023)	$9.70_{0.56}$	$9.83_{0.57}$	$9.99_{0.57}$
ZymPro	$53.17_{1.04}$	$67.22_{0.98}$	$72.44_{0.94}$

4.4 Analysis of Process-Corrected Generation

To answer **RQ3**, Figure 6 compares ZymPro with TSDiff (Kim et al., 2024) and GoFlow (Galustian et al., 2025) for four representative reactions from unseen reaction rules, with all model trajectories initialized at the same reactant geometry. The plot shows two reaction-center distances: the breaking-bond distance on the horizontal axis and the forming-bond distance on the vertical axis. The highest-energy internal MolecularGSM images serve as path-derived TS references. In this projection, ZymPro endpoints are closer to the computational references than either baseline in all four cases. The displayed ZymPro paths generally track the computational reference paths more closely than the baseline paths, with an overall decrease in forming-bond distance. Appendix A.12 provides ten reaction schemes and baseline comparisons; Appendix A.9 details baseline adaptations and geometric evaluation.

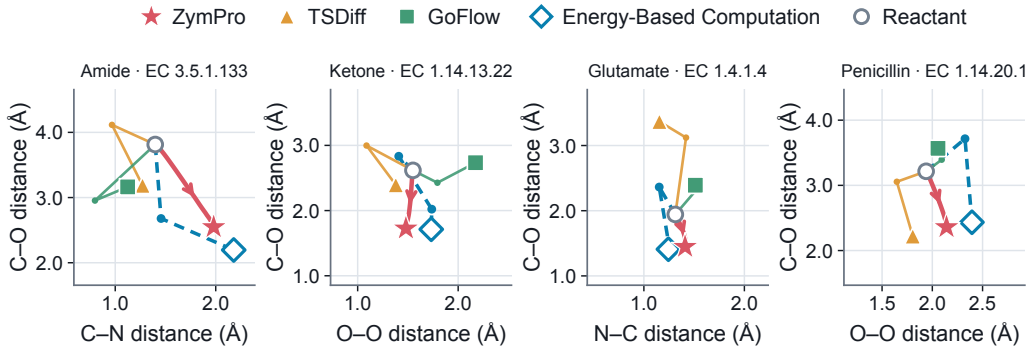

Figure 6: Reaction-center trajectories and TS predictions with shared reactant initialization.

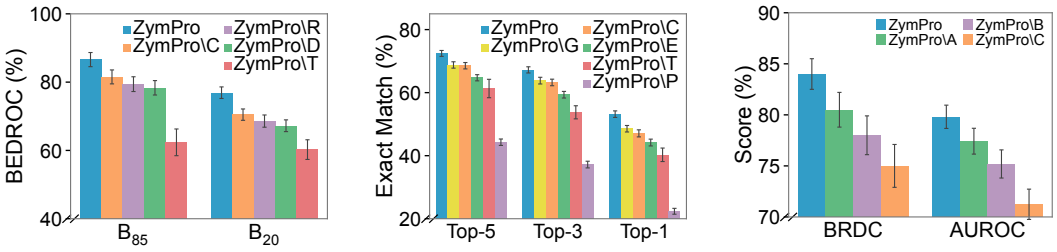

(a) Compatibility (b) Product Prediction (c) Bond Rearrangement

Figure 7: Ablations on enzyme compatibility ranking, product prediction, and bond rearrangement.

4.5 Ablation Studies

We conduct ablation studies to examine the contributions of each component. ZymPro\T replaces transition-state generation with product-endpoint diffusion. ZymPro\G removes TS conditioning while retaining enzyme information. ZymPro\C, ZymPro\A, and ZymPro\B remove the entire Corrector, Advance Correction, and Return Correction. For compatibility ranking, ZymPro\R and ZymPro\D remove transition-state evidence and the direct enzyme–reactant interaction information. Figure 7 summarizes the results. More details are provided in Appendix A.5. We observe that:

- ZymPro\T underperforms the full model on both downstream tasks, with compatibility BEDROC_{85} decreasing from 86.60% to 62.37% and Top-1 from 53.17% to 40.28%. ZymPro\G reduces product Top-1 to 48.55%, supporting its contribution beyond enzyme information.
- ZymPro\R and ZymPro\D reduce compatibility ranking performance, supporting the complementary roles of transition-state and direct enzyme–reactant evidence in enzyme selection.
- ZymPro\C reduces bond forming/breaking direction accuracy (BRDC) from 84.0% to 75.0% and AUROC for identifying changing bonds from 0.7981 to 0.7124, and degrades both downstream tasks, while ZymPro\A and ZymPro\B also lower both metrics. These results support the Corrector's role in bond localization, direction prediction, and downstream performance.

5 Related Work

Transition-State Modeling. Classical transition-state search uses growing-string methods (Peters et al., 2004; Zimmerman, 2013) and lower-cost energies such as GFN2-xTB (Bannwarth et al., 2019). OA-ReactDiff jointly models reactants, transition states, and products (Duan et al., 2023); TSDiff and GoFlow use diffusion or flow matching (Kim et al., 2024; Galustian et al., 2025); React-OT, TS-DFM, SphereDiff-TS, and MolGEN explore transport, distance-based, reaction-center, and joint-pathway formulations (Duan et al., 2025; Luo et al., 2026; Zhao et al., 2026; Tuo et al., 2026). Most use product information; ZymPro generates transition states from enzymes and reactants.

Product Prediction and Enzyme Compatibility. Molecular Transformer, Chemformer, and ReactionT5 predict products (Schwaller et al., 2019; Irwin et al., 2022; Sagawa & Kojima, 2025); Graph2SMILES uses graph-to-sequence decoding (Tu & Coley, 2022); NERF, ReactionSink, and MEGAN redistribute electrons or edit graphs (Bi et al., 2021; Meng et al., 2023; Sacha et al., 2021). ESP, ProSmith, EnzRank, FusionESP, and EZSpecificity model enzyme compatibility (Kroll et al., 2023; 2024; Upadhyay et al., 2023; Du et al., 2025; Cui et al., 2025). CLIPZyme uses pseudo-TS alignment (Mikhael et al., 2024); ReactZyme uses multi-view retrieval (Hua et al., 2024).

6 Conclusion and Future Work

ZymPro combines pairwise enzyme-reactant interactions with energy-based process supervision and an process-informed corrector to generate path-derived transition-state estimates. The generated structures condition product prediction, while TS-supervised features support enzyme compatibility ranking. Experiments under rule- and product-disjoint settings demonstrate improvements over the baselines, and ablations support the contribution of transition-state supervision. Future work will explore higher-level electronic structure calculations with explicit catalytic environments and multiple reaction paths to improve transition-state fidelity and uncertainty estimation. Prospective enzyme screening and experimental validation can assess enzyme selection for desired transformations.

AI USE STATEMENT

Generative AI tools were used to aid and polish writing, assist with literature retrieval and discovery, and help draft portions of the manuscript. All AI-assisted text and retrieved information were reviewed and verified by the authors, who take responsibility for the final content of this work.

REPRODUCIBILITY STATEMENT

Appendix A provides model and training details, data preparation and split procedures, and evaluation protocols for ZymPro and the baselines. Code is available at `https://anonymous.4open.science/r/ZymPro-26E4/`.

REFERENCES

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. URL https://doi.org/10.1021/acs.jctc.8b01176.

Hangrui Bi, Hengyi Wang, Chence Shi, Connor Coley, Jian Tang, and Hongyu Guo. Non-autoregressive electron redistribution modeling for reaction prediction. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 904–913. PMLR, 2021. URL https://proceedings.mlr.press/v139/bi21a.html.

Rebecca Blomberg, Hajo Kries, Daniel M. Pinkas, Peer R. E. Mittl, Markus G. Grütter, Heidi K. Privett, Stephen L. Mayo, and Donald Hilvert. Precision is essential for efficient catalysis in an evolved kemp eliminase. *Nature*, 503:418–421, 2013. doi: 10.1038/nature12623. URL https://doi.org/10.1038/nature12623.

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(8090):639–647, 2025. doi: 10.1038/s41586-025-09697-2. URL https://doi.org/10.1038/s41586-025-09697-2.

Zhenjiao Du, Weimin Fu, Xiaolong Guo, Doina Caragea, and Yonghui Li. FusionESP: Improved enzyme–substrate pair prediction by fusing protein and chemical knowledge. *Journal of Chemical Information and Modeling*, 65(6):2806–2817, 2025. doi: 10.1021/acs.jcim.4c02357. URL https://doi.org/10.1021/acs.jcim.4c02357.

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*, 3:1045–1055, 2023. doi: 10.1038/s43588-023-00563-7. URL https://doi.org/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(4):615–626, 2025. doi: 10.1038/s42256-025-01010-0. URL https://doi.org/10.1038/s42256-025-01010-0.

Stephen D. Fried, Sayan Bagchi, and Steven G. Boxer. Extreme electric fields power catalysis in the active site of ketosteroid isomerase. *Science*, 346(6216):1510–1514, 2014. doi: 10.1126/science.1259802. URL https://doi.org/10.1126/science.1259802.

Leonard Galustian, Konstantin Mark, Johannes Karwounopoulos, Maximilian P.-P. Kovar, and Esther Heid. GoFlow: Efficient transition state geometry prediction with flow matching and E(3)-equivariant neural networks. *Digital Discovery*, 4(12):3492–3501, 2025. doi: 10.1039/D5DD00283D. URL https://doi.org/10.1039/D5DD00283D.

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. URL https://doi.org/10.1039/D3SC02048G.

Chenqing Hua, Bozitao Zhong, Sitao Luan, Liang Hong, Guy Wolf, Doina Precup, and Shuangjia Zheng. ReactZyme: A benchmark for enzyme-reaction prediction. In *Advances in Neural Information Processing Systems*, volume 37, pp. 26415–26442, 2024. doi: 10.52202/079017-0832. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/2e68b2367d2e0bc8dd6f0ff86e07c2eb-Abstract-Datasets_and_Benchmarks_Track.html.

Ross Irwin, Spyridon Dimitriadis, Jiazhen He, and Esben Jannik Bjerrum. Chemformer: A pre-trained transformer for computational chemistry. *Machine Learning: Science and Technology*, 3(1):015022, 2022. doi: 10.1088/2632-2153/ac3ffb. URL https://doi.org/10.1088/2632-2153/ac3ffb.

Xiaohong Ji, Zhen Wang, Zhifeng Gao, Hang Zheng, Linfeng Zhang, Guolin Ke, and Weinan E. Exploring molecular pretraining model at scale. In *Advances in Neural Information Processing Systems*, volume 37, pp. 46956–46978, 2024. doi: 10.52202/079017-1489. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/53923bb44655a7defb31c7744c01b62b-Abstract-Conference.html.

Jisan Kim, Ri Han, and Yoonji Lee. CYP-MAP: A multi-level GNN model for site of metabolism and metabolite prediction. https://github.com/qwon135/CYP-MAP, 2025. Software repository.

Seonghwan Kim, Jeheon Woo, and Woo Youn Kim. Diffusion-based generative AI for exploring transition states from 2D molecular graphs. *Nature Communications*, 15(1): 341, 2024. doi: 10.1038/s41467-023-44629-6. URL https://doi.org/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(1):2787, 2023. doi: 10.1038/s41467-023-38347-2. URL https://doi.org/10.1038/s41467-023-38347-2.

Alexander Kroll, Sahasra Ranjan, and Martin J. Lercher. A multimodal transformer network for protein-small molecule interactions enhances predictions of kinase inhibition and enzyme–substrate relationships. *PLOS Computational Biology*, 20(5):e1012100, 2024. doi: 10.1371/journal.pcbi.1012100. URL https://doi.org/10.1371/journal.pcbi.1012100.

Alpha A. Lee, Qingyi Yang, Vishnu Sresht, Peter Bolgar, Xinjun Hou, Jacquelyn L. Klug-McLeod, and Christopher R. Butler. Molecular transformer unifies reaction prediction and retrosynthesis across pharma chemical space. *Chemical Communications*, 55:12152–12155, 2019. doi: 10.1039/C9CC05122H. URL https://doi.org/10.1039/C9CC05122H.

Chenao Li, Shuo Yan, and Enyan Dai. UniZyme: A unified protein cleavage site predictor enhanced with enzyme active-site knowledge. In *Advances in Neural Information Processing Systems*, volume 38, pp. 18507–18535, 2025. doi: 10.52202/085713-0624. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/1ac3030fc57850b0fb11dfe9d4880ad7-Abstract-Conference.html.

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

Ziqiao Meng, Peilin Zhao, Yang Yu, and Irwin King. Doubly stochastic graph-based non-autoregressive reaction prediction. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pp. 4064–4072, 2023. doi: 10.24963/ijcai.2023/452. URL https://doi.org/10.24963/ijcai.2023/452.

Peter Mikhael, Itamar Chinn, and Regina Barzilay. CLIPZyme: Reaction-conditioned virtual screening of enzymes. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 35647–35663. PMLR, 2024. URL https://proceedings.mlr.press/v235/mikhael24a.html.

Alex Morehead and Jianlin Cheng. FlowDock: Geometric flow matching for generative protein-ligand docking and affinity prediction. *Bioinformatics*, 41(Supplement_1):i198–i206, 2025. doi: 10.1093/bioinformatics/btaf187. URL https://doi.org/10.1093/bioinformatics/btaf187.

Masao Ohashi, Cooper S. Jamieson, Yujuan Cai, Dan Tan, Daiki Kanayama, Man-Cheng Tang, Sarah M. Anthony, Jason V. Chari, Joyann S. Barber, Elias Picazo, Thomas B. Kakule, Shugeng Cao, Neil K. Garg, Jiahai Zhou, K. N. Houk, and Yi Tang. An enzymatic alder–ene reaction. *Nature*, 586(7827):64–69, 2020. doi: 10.1038/s41586-020-2743-5. URL https://doi.org/10.1038/s41586-020-2743-5.

Baron Peters, Andreas Heyden, Alexis T. Bell, and Arup K. Chakraborty. A growing string method for determining transition states: Comparison to the nudged elastic band and string methods. *The Journal of Chemical Physics*, 120(17):7877–7886, 2004. doi: 10.1063/1.1691018. URL https://doi.org/10.1063/1.1691018.

Mikołaj Sacha, Mikołaj Błaż, Piotr Byrski, Paweł Dąbrowski-Tumański, Mikołaj Chromiński, Rafał Loska, Paweł Włodarczyk-Pruszyński, and Stanisław Jastrzębski. Molecule edit graph attention network: Modeling chemical reactions as sequences of graph edits. *Journal of Chemical Information and Modeling*, 61(7):3273–3284, 2021. doi: 10.1021/acs.jcim.1c00537. URL https://doi.org/10.1021/acs.jcim.1c00537.

Tatsuya Sagawa and Ryosuke Kojima. ReactionT5: A pre-trained transformer model for accurate chemical reaction prediction with limited data. *Journal of Cheminformatics*, 17:126, 2025. doi: 10.1186/s13321-025-01075-4. URL https://doi.org/10.1186/s13321-025-01075-4.

Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations*, 2022. URL https://arxiv.org/abs/2202.00512.

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

Philippe Schwaller, Teodoro Laino, Theophile 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. doi: 10.1021/acscentsci.9b00576. URL https://doi.org/10.1021/acscentsci.9b00576.

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. URL https://arxiv.org/abs/2010.02502.

Siyang Tian, Xuan Cao, Russell Greiner, Carin Li, Anchi Guo, and David S. Wishart. CyProduct: A software tool for accurately predicting the byproducts of human cytochrome P450 metabolism. *Journal of Chemical Information and Modeling*, 61(6):3128–3140, 2021. doi: 10.1021/acs.jcim.1c00144. URL https://doi.org/10.1021/acs.jcim.1c00144.

Jean-François Truchon and Christopher I. Bayly. Evaluating virtual screening methods: Good and bad metrics for the "early recognition" problem. *Journal of Chemical Information and Modeling*, 47(2):488–508, 2007. doi: 10.1021/ci600426e. URL https://doi.org/10.1021/ci600426e.

Zhengkai Tu and Connor W. Coley. Permutation invariant graph-to-sequence model for template-free retrosynthesis and reaction prediction. *Journal of Chemical Information and Modeling*, 62(15):3503–3513, 2022. doi: 10.1021/acs.jcim.2c00321. URL https://doi.org/10.1021/acs.jcim.2c00321.

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

Vikas Upadhyay, Veda Sheersh Boorla, and Costas D. Maranas. Rank-ordering of known enzymes as starting points for re-engineering novel substrate activity using a convolutional neural network. *Metabolic Engineering*, 78:171–182, 2023. doi: 10.1016/j.ymben.2023.06.001. URL https://doi.org/10.1016/j.ymben.2023.06.001.

Yitian Wang, Jingxin Rao, Wei Zhang, Yuqi Shi, Chuanlong Zeng, Rongrong Cui, Yinquan Wang, Jiacheng Xiong, Xutong Li, and Mingyue Zheng. MetaReact: A reaction-aware transformer for end-to-end prediction of drug metabolism. *bioRxiv*, 2026. doi: 10.64898/2026.03.14.711529. URL https://doi.org/10.64898/2026.03.14.711529.

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. doi: 10.1021/cr0503106. URL https://doi.org/10.1021/cr0503106.

Chong Zhao, Pan Li, Shu Zhang, Chuanhao Li, Zhaopeng Li, Yixin Tang, Keyan Linghu, Lei Tang, and Yuanyong Yang. SphereDiff-TS: Sphere space diffusion modeling for accurate 3D transition state geometry prediction. *ACS Omega*, 11(17):26048–26059, 2026. doi: 10.1021/acsomega.6c03301. URL https://doi.org/10.1021/acsomega.6c03301.

Paul M. Zimmerman. Reliable transition state searches integrated with the growing string method. *Journal of Chemical Theory and Computation*, 9(7):3043–3050, 2013. doi: 10.1021/ct400319w. URL https://doi.org/10.1021/ct400319w.

A Additional Methodological and Experimental Details

This appendix describes the enzyme-swap diagnostic, offline construction of transition-state supervision, generator training and sampling, downstream interfaces, and evaluation protocols. Section A.5 separates downstream implementation from the supplementary ablation results. Notation follows the main text.

A.1 Enzyme-Swap Diagnostics

Data and Candidates. We rebuild labels from the 102 substrates in the released CYP-MAP test set (Kim et al., 2025), obtaining 222 canonical reactant-product groups over 90 target-bearing reactants. RDKit canonicalization retains stereochemistry, formal charge, isotopes, and all product components. Each group uses the same nine candidates: CYP1A2, CYP2A6, CYP2B6, CYP2C8, CYP2C9, CYP2C19, CYP2D6, CYP2E1, and CYP3A4. All recorded isoforms for that group are positives; the others are unannotated alternatives. Candidate construction uses only these labels and retains every group, giving 1,998 conditions and 2,922 positive-alternative pairs per model. We additionally evaluate ZymPro on the same diagnostic dataset, which is excluded from its training data, and report its results in the tables below.

Frozen Target Scoring. MetaReact (Wang et al., 2026) and EC3-MT use length-normalized, teacher-forced target log-likelihood, with strict ReactSeq round-trip validation for MetaReact. CYP-MAP (Kim et al., 2025) uses its official five-checkpoint ensemble, and CyProduct (Tian et al., 2021) uses the unmodified official program; both retain a native score only for an exact complete-target match. EC3-MT uses the fixed step-7,000 checkpoint for Table 3. All parameters remain frozen, and scores are compared only within each model.

Ranking and Uncertainty. Available target scores rank above missing scores, which form one bottom tie block. Top-1 is the positive fraction of the highest tie block; MRR averages the reciprocal rank of the first positive under uniform ordering within ties. All 222 groups are retained. The

random-guess Top-1 reference is 26.58%, averaged from each group's positive count divided by nine. Figure 2 uses 95% percentile intervals from 1,000 bootstrap samples clustered by canonical reactant. Table 4 separates score availability from score-only comparisons, for which both candidates must have numeric scores. Score-only preference is the annotated enzyme's win rate, with half credit for ties. The tie rate is the fraction of all 2,922 positive-alternative pairs with equal stored numeric scores or with both scores missing. Equality uses the stored values without a tuned tolerance. Coverage uses the 1,998 candidate conditions as its denominator. Bootstrap resampling retains all groups and pairs from each sampled reactant; pairwise statistics retain pair weighting and Top-1 and MRR retain group weighting.

Table 4: Additional fixed-target results. Coverage is target-score availability. Top-1 and score-only preference include 95% CIs; score-only comparisons use the listed eligible pairs (2,922 total).

Model	Coverage (%)	Top-1 (%)	MRR	Ties (%)	Score-only (%)	Pairs
MetaReact	28.83	37.99 [32.98, 43.82]	0.571	74.54	74.19 [64.24, 82.94]	744
CYP-MAP	12.31	39.44 [33.49, 45.56]	0.569	78.03	77.57 [68.95, 85.53]	272
CyProduct	22.57	44.84 [38.48, 51.41]	0.614	55.41	64.00 [54.38, 73.36]	418
EC3-MT	100.00	24.21 [19.49, 30.35]	0.465	81.72	51.68 [50.14, 53.21]	2,922
ZymPro	100.00	78.19 [72.81, 84.57]	0.860	1.47	86.34 [81.12, 91.63]	2,922

Generated-Product Invariance. Generation uses only the reactant and enzyme. MetaReact aggregates its official ten enumerations and ten beams by inverse-squared-rank votes, retaining complete products and strict source/product ReactSeq round-trip checks. CYP-MAP uses native rank 1; CyProduct uses the maximum native score. EC3-MT uses beam size 5, maximum length 224, and no length penalty, retaining its first beam. Unchanged product is the fraction of all 2,922 swaps yielding identical valid top-ranked complete-product sets. Native top-rank ties are compared as sets; missing or failed outputs contribute zero. Table 5 separates this rate from invariance conditional on two valid outputs. Its intervals use 1,000 bootstrap replicates clustered by canonical reactant.

Table 5: Product invariance after enzyme replacement. Each model uses 2,922 swaps. Conditional rates use only the listed valid pairs (95% CIs).

Model	Valid pairs	Unchanged (%)	Conditional (%)
MetaReact	712	23.34	95.79 [88.92, 100.00]
CYP-MAP	1,004	31.72	92.33 [86.16, 97.36]
CyProduct	1,077	23.03	62.49 [53.03, 71.64]
EC3-MT	2,256	69.37	89.85 [86.89, 92.63]
ZymPro	2,801	11.35	11.84 [8.83, 17.58]

A.2 Transition-State Target Construction

We use "TS estimate" to denote the substrate geometry associated with the highest-energy internal image of the enzyme-conditioned GSM path, which serves as the model's prediction target. This path-derived target is our *operational transition-state reference*. The ordered path supplies the return and advance correction targets in Section 3.1.

Offline Endpoint Preparation. For each eligible reaction, we dock the reactant and product endpoints separately into the enzyme pocket with FlowDock (Morehead & Cheng, 2025). We use the released `esmfold_prior_paper_weights-EMA` checkpoint, 20 inference steps, one sample per endpoint, and the rank-1 pose. The endpoints are repaired when needed, placed in a common atom order, optimized under the same constraints for at most 150 steps, and connected by a double-ended molecularGSM path (Peters et al., 2004; Zimmerman, 2013) using GFN2-xTB energies and gradients (Bannwarth et al., 2019). Each path uses nine nodes, 80 maximum optimization iterations, 30 step-optimization iterations, mobile-coordinate RMS gradient tolerance $5\times10^{-4}\ E_h/\mathrm{bohr}$, node-addition tolerance 0.1, SSM displacement limit 0.8, and minimum spacing 5.0. The molecularGSM product limit is 100 and its final-TS type is 1. The total charge and number of unpaired electrons are assigned to the complete substrate–pocket cluster. We remove failed path searches, atom-correspondence errors, and structurally invalid outputs.

Energy-Evaluation System. The pocket includes annotated active-site residues and residues within 5.0 Å of substrate heavy atoms in either docked endpoint, retaining essential catalytic waters, cofactors, and first-shell metal ligands. The endpoint selections are united in a common protein frame, with one atom set and mapping throughout the path. Preparation uses reference pH 7.0; His protonation follows local hydrogen bonds, with catalytic and metal-binding residues assigned individually. Peripheral truncations are hydrogen-capped along the original bonds (C–H: 1.09 Å), retaining functional backbone segments. Boundary Cα and peripheral backbone heavy atoms are fixed; substrates, catalytic side chains, relevant hydrogens, and catalytic waters remain mobile. Charge is determined from the complete prepared cluster; closed-shell systems use zero unpaired electrons, with metal and radical states specified individually. GFN2-xTB evaluates the entire cluster, including fixed atoms. Pocket coordinates are implicit in $E_{\mathrm{xTB}}(\mathbf{X})$; exported targets retain only mapped reacting-molecule atoms. Prepared clusters, fixed-atom masks, protonation, capping, and charge/spin assignments are retained.

Transition-State Target. For reaction i, molecularGSM produces the ordered path $\Gamma_i = (\mathbf{X}_i^{(0)}, \ldots, \mathbf{X}_i^{(K)})$ defined in Eq. (3). For total cluster charge q_i^{chg} and unpaired-electron count n_i^{unp}, GFN2-xTB evaluates each frame as $E_{i,k} = E_{\mathrm{xTB}}(\mathbf{X}_i^{(k)}; q_i^{\mathrm{chg}}, n_i^{\mathrm{unp}})$ and supplies the gradient used by Eq. (2). After convergence, we select the highest-energy internal frame:

$$\begin{aligned} k_i^\star &= \arg\max_{1\le k\le K-1} E_{i,k}, \\ \mathbf{X}_i^{TS} = \mathbf{X}_i^{(k_i^\star)}, \qquad D_{i,ab}^{TS} &= \left\| \mathbf{x}_{i,a}^{(k_i^\star)} - \mathbf{x}_{i,b}^{(k_i^\star)} \right\|_2 . \end{aligned} \tag{13}$$

The substrate geometry of this semiempirical path maximum defines the operational transition-state reference used for supervision and evaluation. Endpoint docking, path frames, energies, and gradients belong to the offline supervision pipeline. The learned model encodes the enzyme and reactants separately; neither training nor inference takes their docked complex as an input.

A.3 Transition-State Generator Details

This subsection expands the three learned components in Section 3.1: interaction memory, residual denoising, and reaction path correction. It concludes with their joint objective.

In Figure 3, the Denoiser is the Atom–Pair Transformer f_ϕ, and Sampling Update denotes the deterministic DDIM step. The molecular views schematically illustrate distances reconstructed as $\mathbf{D}^R + s_\Delta \mathbf{R}_t$; $\mathbf{R}_t$ itself is a distance residual, rather than an atomic coordinate configuration.

Input Geometry and Cross-System Features. Let $D_{RE} \in \mathbb{R}^{n\times l}$ denote the atom-residue distance channel. We set $D_{RE} = 0$ throughout training and inference, so the interaction encoder does not require an enzyme-reactant docking pose. This channel is distinct from the within-reactant distance matrix $D_{ab}^R = \|x_a - x_b\|_2$. Enzyme conditioning enters through residue features and learned cross-pair interactions.

Pairwise Interaction Encoding. The reactant pair state includes the within-reactant geometry:

$$B_{RR} \leftarrow \mathrm{Sym}\big(B_{RR} + \mathrm{MLP}_D\big(\mathrm{RBF}\big(\mathrm{clip}(D^R, 0, 40\,\text{Å})\big)\big)\big), \tag{14}$$

where RBF is a fixed 32-center Gaussian basis spanning $[0, 40]$ Å, MLP_D is a learned Linear-GELU-Linear projection, and Sym symmetrizes the pair state. The atom and residue state vectors h_a^R and h_r^E are projected to the pair width and added to form $C_{ar} = W_R h_a^R + W_E h_r^E$, where W_R and W_E are learned linear maps into the common cross-pair feature space. The first axial attention operates over atoms while holding residue r fixed. For head h, its attention logit is

$$\xi_{ab|r}^h = \frac{(q_{ar}^h)^\top k_{br}^h}{\sqrt{d_h}} + \mathrm{clip}\big((w_{RR}^h)^\top B_{RR,ab}, -b_{\max}, b_{\max}\big). \tag{15}$$

The queries, keys, and values are learned projections of normalized cross-pair states. The second axial attention operates over enzyme residues while holding atom a fixed:

$$\xi_{rs|a}^h = \frac{(\widetilde{q}_{ar}^h)^\top \widetilde{k}_{as}^h}{\sqrt{d_h}} + \mathrm{clip}\big((w_{EE}^h)^\top B_{EE,rs}, -b_{\max}, b_{\max}\big). \tag{16}$$

We set $b_{\max} = 8$. Each attention output is combined through a $\tanh$-gated residual update and layer normalization. The two axial updates produce the cross-pair state B_{RE} in the shared interaction memory $\mathcal{M}(\mathcal{E}, \mathcal{R})$ described in Section 3.1. Padding masks exclude invalid atoms and residues, and $B_{ER} = B_{RE}^{\top}$. Thus, each cross-pair state depends on its atom and residue content, while B_{RR} and B_{EE} organize information along the two axes. Setting $D_{RE} = 0$ does not set B_{RE} to zero: the latter is the learned cross-pair state produced by these axial attention operations.

Reactant-Relative Residual Diffusion. Following Eq. (4), the clean target is $\mathbf{R}_0 = (\mathbf{D}^{TS} - \mathbf{D}^{R})/s_{\Delta}$, and $\mathbf{R}_t$ is its noisy diffusion state. The denoiser f_ϕ initializes

$$\begin{aligned} H_{R,t}^{(0)} &= H_R, \\ B_t^{(0)} &= \mathrm{PairMask}(\mathrm{Emb}_R(\mathbf{R}_t) + W_B B_{RR} + \mathrm{Emb}_{\mathrm{sc}}(\mathbf{R}_{\mathrm{sc}})), \end{aligned} \tag{17}$$

where Emb_R embeds the noisy scalar residual, W_B projects the within-reactant pair state to the dynamic pair width, $\mathrm{Emb}_{\mathrm{sc}}$ embeds the optional clean-residual self-condition $\mathbf{R}_{\mathrm{sc}}$, and $\mathrm{PairMask}$ symmetrizes the pair state and removes invalid atom pairs while retaining diagonal pair features. A missing self-condition is represented by $\mathbf{R}_{\mathrm{sc}} = 0$. The timestep embedding is $e_t = \mathrm{MLP}_t(\mathrm{TimeEmb}(t))$. For head h in denoising block ℓ, atom self-attention and cross-attention from atoms to enzyme residues use

$$\begin{aligned} \xi_{ab}^{RR,\ell,h} &= \frac{(q_a^{\ell,h})^{\top} k_b^{\ell,h}}{\sqrt{d_h}} + (w_{RR}^{\ell,h})^{\top}(B_{RR,ab} + B_{t,ab}^{(\ell-1)}), \\ \xi_{ar}^{RE,\ell,h} &= \frac{(\widetilde{q}_a^{\ell,h})^{\top} k_r^{E,\ell,h}}{\sqrt{d_h}} + (w_{RE}^{\ell,h})^{\top} B_{RE,ar}. \end{aligned} \tag{18}$$

Here $B_t^{(\ell-1)}$ is the dynamic reactant pair state entering block ℓ, d_h is the attention-head width, and all Q/K/V and pair projections are learned. The cross-pair state $B_{RE,ar}$ depends on both the atom and residue and therefore provides a pair-specific attention bias. Candidate identity also enters through the residue keys and values, so attention from atoms to enzyme residues combines residue content with the structured cross-pair context.

We summarize the time-conditioned atom updates as

$$\widetilde{H}_{R,t}^{(\ell)} = \mathrm{AtomBlock}_{\ell}\left(H_{R,t}^{(\ell-1)}, H_E; B_{RR} + B_t^{(\ell-1)}, B_{RE}, e_t\right). \tag{19}$$

Here, $\mathrm{AtomBlock}_{\ell}$ applies pair-biased atom self-attention, cross-attention from atoms to residues, and a feed-forward update; adaptive layer normalization uses e_t in all three operations. The updated atom states refresh the reactant pair representation through a projected multiplicative update:

$$O_{ab}^{(\ell)} = W_{\mathrm{opm},O}\left[(W_{\mathrm{opm},L}\widetilde{H}_{R,t,a}^{(\ell)}) \odot (W_{\mathrm{opm},R}\widetilde{H}_{R,t,b}^{(\ell)})\right]. \tag{20}$$

The cross-pair states provide a second pair update. We first average over valid enzyme residues and then combine the resulting atom-wise contexts. Let $m_r^E \in \{0,1\}$ indicate whether residue r is valid:

$$\begin{aligned} \overline{B}_{RE,a} &= \frac{\sum_r m_r^E B_{RE,ar}}{\max(1, \sum_r m_r^E)}, \\ P_{ab}^{RE} &= \mathrm{MLP}_P\big(\mathrm{CONCAT}(\overline{B}_{RE,a}, \overline{B}_{RE,b})\big). \end{aligned} \tag{21}$$

The pair-state update is

$$\begin{aligned} \widetilde{B}_{t,ab}^{(\ell)} &= B_{t,ab}^{(\ell-1)} + \tanh(\gamma_{\mathrm{mul}}) O_{ab}^{(\ell)} + \tanh(\gamma_{\mathrm{cross}}) P_{ab}^{RE}, \\ \overline{B}_t^{(\ell)} &= \mathrm{Sym}\left(\widetilde{B}_t^{(\ell)}\right), \\ B_{t,\mathrm{tri}}^{(\ell)} &= \mathrm{Sym}\Big(\mathrm{Tri}_{\ell}(\overline{B}_t^{(\ell)}; \mathbf{m}_{RR})\Big), \\ H_{R,t}^{(\ell)} &= \widetilde{H}_{R,t}^{(\ell)} + \tanh(\gamma_{\mathrm{pair}})\,\mathrm{PairToAtom}\left(B_{t,\mathrm{tri}}^{(\ell)}; \mathbf{m}_{RR}\right), \\ B_t^{(\ell)} &= \mathbf{m}_{RR} \odot B_{t,\mathrm{tri}}^{(\ell)}, \end{aligned} \tag{22}$$

where the three γ's are learned scalar gates and Tri_{ℓ} applies triangle multiplication, triangle attention, and a pair transition while masking their residual updates. $\mathbf{m}_{RR}$ is the valid reactant-pair mask.

Aggregation from pairs to atoms uses the same mask before invalid pair entries are removed. After $L = 2$ blocks, the base denoising outputs are

$$\begin{aligned} \widetilde{\mathbf{R}}_0^{(t)} &= \mathbf{m}_{RR} \odot \mathrm{SymZeroDiag}\Big(\mathrm{Head}_R\Big(B_t^{(L)}\Big)\Big), \\ H_{R,t}^{\mathrm{out}} &= H_{R,t}^{(L)}, \qquad B_{RR,t}^{\mathrm{out}} = B_t^{(L)}. \end{aligned} \tag{23}$$

Here $\mathrm{SymZeroDiag}$ symmetrizes the scalar residual, sets its diagonal to zero, and the pair mask removes invalid atom pairs. The same mask is used by the distance losses. Equation (5) then applies the learned path correction to obtain $\widehat{\mathbf{R}}_0^{(t)}$.

Reaction Path Correction. For each denoiser estimate $\widetilde{\mathbf{R}}_0^{(t)}$, we interpolate aligned adjacent reaction frames and select the nearest path position $\mathbf{R}_\Gamma^{(j)}$ by valid-pair distance error. Advancement follows the cached structural arc toward the path maximum by at most 1 Å, stopping at the maximum. The return and advance targets in Eq. (7) are

$$\mathbf{C}_r^\star(\widetilde{\mathbf{R}}_0^{(t)}) = \mathbf{R}_\Gamma^{(j)} - \widetilde{\mathbf{R}}_0^{(t)}, \qquad \mathbf{C}_a^\star(\widetilde{\mathbf{R}}_0^{(t)}) = \mathbf{R}_\Gamma^{(j+1)} - \mathbf{R}_\Gamma^{(j)}.$$

When j is the terminal transition-state position, we set $\mathbf{R}_\Gamma^{(j+1)} = \mathbf{R}_\Gamma^{(j)} = \mathbf{R}_0$, yielding zero advance correction. Both targets vanish when the estimate reaches the path maximum $\mathbf{R}_0$. The nearest-path RMS is retained only as a diagnostic of path assignment and is not itself a supervision target. Target construction is detached from gradient computation.

We supervise the Corrector on the current denoiser estimate and on a path frame perturbed by 0.15 Å coordinate noise. Given the estimate, diffusion time t, and interaction memory $\mathcal{M}$, the Corrector g_ϕ first predicts pre-bounded branches $\overline{\mathbf{C}}_\phi^r$ and $\overline{\mathbf{C}}_\phi^a$. A shared factor controls their total amplitude before returning the corrections $\mathbf{C}_\phi^r$ and $\mathbf{C}_\phi^a$ used in Eq. (5). With $\overline{\mathbf{C}}_\phi = \overline{\mathbf{C}}_\phi^r + \overline{\mathbf{C}}_\phi^a$,

$$\begin{aligned} \kappa_\phi(t, \overline{\mathbf{C}}_\phi) &= 0.25 \tanh[\gamma + 0.10 \tanh(\mathrm{MLP}_s(e_t))] \min\left(1, \frac{\rho}{\max\{\mathrm{RMS}_{\mathbf{m}_{RR}}(\overline{\mathbf{C}}_\phi), \varepsilon\}}\right), \\ \mathbf{C}_\phi^r &= \kappa_\phi(t, \overline{\mathbf{C}}_\phi)\overline{\mathbf{C}}_\phi^r, \qquad \mathbf{C}_\phi^a = \kappa_\phi(t, \overline{\mathbf{C}}_\phi)\overline{\mathbf{C}}_\phi^a. \end{aligned}$$

The outputs of g_ϕ in Eq. (5) are the bounded fields $\mathbf{C}_\phi^r$ and $\mathbf{C}_\phi^a$; their sum is added directly to $\widetilde{\mathbf{R}}_0^{(t)}$. Here $\mathrm{RMS}_{\mathbf{m}_{RR}}$ is computed under the valid-pair mask, ε is a numerical stabilizer, and e_t is the diffusion-time embedding. The Corrector uses the same residual-target scale s_Δ as the denoiser and correction losses. The formal configuration uses a bounded correction gate with maximum absolute value 0.25 and a time-offset scale of 0.10. The RMS bound is $\rho = 1\,\text{Å}/s_\Delta$, applied to the summed raw correction before gating; this bound is separate from the path-advancement step above. The initialization $\gamma = \mathrm{atanh}(0.1)$ gives an initial gate scale near 0.025.

The return and advance losses in Eq. (8) are implemented as $\mathcal{L}_c$, evaluated on both the current and perturbed estimates to supervise the two applied corrections:

$$\mathcal{L}_c = \sum_{\chi \in \{\mathrm{live}, \mathrm{pert}\}} \left[\mathrm{SmoothL1}_{0.1}(\mathbf{C}_{\phi,\chi}^r, \mathbf{C}_{r,\chi}^\star) + \mathrm{SmoothL1}_{0.1}(\mathbf{C}_{\phi,\chi}^a, \mathbf{C}_{a,\chi}^\star)\right]. \tag{24}$$

Here live denotes the current denoiser estimate and pert the perturbed path estimate. $\mathrm{SmoothL1}_{0.1}$ instantiates the robust regression loss ℓ in Eq. (8); it is quadratic near zero and linear beyond an absolute error of 0.1.

The auxiliary correction terms set both fields to zero at $\mathbf{R}_0$, align their direction and magnitude with the local path, and regularize the corrected distance matrix. We denote these terms by $\mathcal{L}_{\mathrm{fix}}$, $\mathcal{L}_{\mathrm{profile}}$, and $\mathcal{L}_{\mathrm{valid}}$, with $\mathcal{L}_{\mathrm{fix}} = \mathrm{SmoothL1}_{0.1}(\mathbf{C}_\phi^r(\mathbf{R}_0), 0) + \mathrm{SmoothL1}_{0.1}(\mathbf{C}_\phi^a(\mathbf{R}_0), 0)$. The profile term covers path progress, distance and energy gap to the peak, local energy slope, peak sharpness and confidence, and the direction and magnitude of the advance correction. Progress is the clipped frame offset from the peak; peak distance is normalized by the corresponding half-path arc length. Energy gap and sharpness use the internal-frame energy range, while confidence combines normalized peak separation and slope consistency. Slope has three classes; advance targets encode RMS amplitude and normalized direction. Scalar targets use SmoothL1 regression, slope uses cross-entropy, and direction uses cosine loss. The corresponding weights are $0.05, 0.05, 0.025, 0.025, 0.01, 0.01, 0.025,$

and 0.025. The validity term penalizes gate deviation from 0.025 and the normalized error between predicted and reference distance-derived Gram matrices, weighted by 0.001 and 0.05.

Complete Training Objective. We instantiate Eq. (8) in four parts: masks and pair weights, denoising losses, auxiliary constraints, and the joint objective.

Masks and pair weights. Before scaling and residual construction, the transition-state target distances are clipped elementwise to $[0, 40]$ Å; the reactant distance matrix D^R in the residual remains unclipped. Distances are internally scaled as $\bar{D} = D/s_D$, where $s_D = 10\,\text{Å}$, while $\mathbf{R}_0 = (\mathbf{D}^{TS} - \mathbf{D}^R)/s_\Delta$ uses $s_\Delta = 2\,\text{Å}$. The one-step distance estimate and its scaled form are

$$\widehat{D}_t^{TS} = D^R + s_\Delta \widehat{\mathbf{R}}_0^{(t)}, \qquad \widehat{\bar{D}}_t^{TS} = \widehat{D}_t^{TS}/s_D.$$

The corresponding noise estimate is

$$\widehat{\epsilon}_t = \frac{\mathbf{R}_t - \sqrt{\bar{\alpha}_t}\widehat{\mathbf{R}}_0^{(t)}}{\sqrt{1-\bar{\alpha}_t}}.$$

For sample i, let m_{ia} be the valid atom mask and h_i indicate that a retained transition-state target is available. The target mask is

$$m_{iab}^{TS} = h_i\, m_{ia} m_{ib}\, \mathbf{1}[a \neq b]. \tag{25}$$

Distance losses are evaluated only on positive enzyme-reaction records with retained transition-state targets. Samples failing reaction path search, atom correspondence, or structural validity are excluded. In the reducer below, $\omega_i \geq 0$ denotes the fixed sample weight attached to a retained record; it rescales optimization only.

Pair weights emphasize short-range and endpoint-relevant relations. Before normalization, the implemented weight is

$$\begin{aligned} w_{iab} = \max\{1,\ & 4\mathbf{1}[D_{iab}^{TS} \leq 4],\ 2\mathbf{1}[D_{iab}^{TS} \leq 8], \\ & 2\mathbf{1}[D_{iab}^{R} \leq 4],\ 1.5\mathbf{1}[D_{iab}^{R} \leq 8], \\ & 4\mathbf{1}[(a,b) \in \mathcal{B}_i^R],\ 6\mathbf{1}[a \in \mathcal{A}_i^{\text{chg}} \ \vee\ b \in \mathcal{A}_i^{\text{chg}}]\}, \end{aligned} \tag{26}$$

where distances are in Å, $\mathcal{B}_i^R$ is the reactant bond set, and $\mathcal{A}_i^{\text{chg}}$ contains atoms involved in endpoint bond or charge changes. These quantities are used only as training-loss weights and are never model inputs. We normalize w_{iab} to unit mean over valid pairs within each molecule, obtaining $\widetilde{w}_{iab}$.

For a pairwise value A_{iab}, mask ν_{iab}, pair weight v_{iab}, and sample weight c_i, let $\mathcal{I}_\nu = \{i : \sum_{a,b} \nu_{iab} v_{iab} > 0\}$. We use the following weighted average throughout the objective:

$$\text{Avg}(A; \nu, v, c) = \frac{1}{\sum_{i \in \mathcal{I}_\nu} c_i} \sum_{i \in \mathcal{I}_\nu} c_i \frac{\sum_{a,b} \nu_{iab} v_{iab} A_{iab}}{\sum_{a,b} \nu_{iab} v_{iab}}. \tag{27}$$

All denominators are lower-bounded numerically. During training, the zero-based schedule index τ is sampled uniformly from the rounded 25-point DDIM grid and jittered by an integer in $\{-3, \ldots, 3\}$. This draw is replaced by a uniform index in $\{0, \ldots, 999\}$ with probability 0.05, followed by a uniform high-noise index in $\{900, \ldots, 999\}$ with probability 0.25. In the mathematical forward-process notation of Eq. (4), this index corresponds to $t = \tau + 1$.

Denoising losses. The reconstruction, residual, and diffusion-noise terms are

$$\begin{aligned} \mathcal{L}_{\text{rec}} &= \text{Avg}\Big((\widehat{\bar{D}}_t^{TS} - \bar{D}^{TS})^2; m^{TS}, \widetilde{w}, \omega \Big), \\ \mathcal{L}_{\text{res}} &= \text{Avg}\Big((\widehat{\mathbf{R}}_0^{(t)} - \mathbf{R}_0)^2; m^{TS}, \widetilde{w}, \omega \Big), \\ \mathcal{L}_{\epsilon} &= \text{Avg}\big((\widehat{\epsilon}_t - \epsilon)^2; m^{TS}, 1, \omega\eta_t \big), \\ \eta_t &= \frac{\min(\text{SNR}_t, 5)}{\text{SNR}_t}, \qquad \text{SNR}_t = \frac{\bar{\alpha}_t}{1-\bar{\alpha}_t}. \end{aligned} \tag{28}$$

With the same masks and weights, the two reconstruction terms satisfy

$$\mathcal{L}_{\text{rec}} = \left(\frac{s_\Delta}{s_D}\right)^2 \mathcal{L}_{\text{res}} = 0.04\, \mathcal{L}_{\text{res}}.$$

They express the same reconstruction error in distance and residual units; their coefficients in the joint objective are retained below.

Auxiliary constraints. Let $e_{iab} = |\widehat{D}^{TS}_{t,iab} - D^{TS}_{iab}|$ in Å, and define

$$\begin{aligned} q_{iab} &= \frac{1}{4} \sum_{\tau \in \{0.5,1,2,4\}} \sigma(\tau - e_{iab}), \\ m^{\mathrm{loc}}_{iab} &= m^{TS}_{iab} \mathbf{1}[D^{TS}_{iab} < 15], \\ \mathcal{L}_{\mathrm{lddt}} &= 1 - \mathrm{Avg}(q; m^{\mathrm{loc}}, 1, \omega). \end{aligned} \tag{29}$$

Here σ is the logistic sigmoid, with numerical distances expressed in Å. Equation (29) is a soft-lDDT surrogate: it replaces hard distance-error thresholds with sigmoid terms. It is used without rescaling its zero-error value, $q_{\max} = \frac{1}{4}[\sigma(0.5) + \sigma(1) + \sigma(2) + \sigma(4)]$, to one. Let $\bar{q}_i$ be the masked mean of q_{iab} over m^{loc}_{iab} for sample i. Its detached value supervises a scalar head for this same surrogate, rather than a discrete lDDT score:

$$\widehat{q}_i = \sigma\left(\mathrm{Head}_q \left(\frac{\sum_a m_{ia} H^{\mathrm{out}}_{R,t,ia}}{\max(1, \sum_a m_{ia})} \right) \right). \tag{30}$$

Let $X^R_i := \mathbf{x}_i$ denote the reactant coordinates from the problem definition and X^{TS}_i the target coordinates derived from the path search. At diffusion step t, a coordinate head predicts

$$\widehat{X}^{TS}_{t,i} = X^R_i + s_x \, \mathrm{Head}_x \left(H^{\mathrm{out}}_{R,t,i} \right), \tag{31}$$

where $s_x = 2\,\text{Å}$ is a fixed coordinate-residual scale. The auxiliary losses are

$$\begin{aligned} \mathcal{L}_q &= \frac{\sum_i \omega_i (\widehat{q}_i - \mathrm{stopgrad}(\bar{q}_i))^2}{\sum_i \omega_i}, \\ \mathcal{L}_{\mathrm{coord}} &= \frac{1}{\sum_{i \in \mathcal{I}_x} \omega_i} \sum_{i \in \mathcal{I}_x} \omega_i \frac{\sum_a m^x_{ia} \| X^{TS,\mathrm{ctr}}_{ia} - \widehat{X}^{TS,\mathrm{ctr}}_{t,ia} \|_2^2}{\sum_a m^x_{ia}}, \\ \mathcal{L}_{\mathrm{near}} &= \mathrm{Avg}\left((\widehat{\bar{D}}^{TS}_t - \bar{D}^{TS})^2; m^{\mathrm{near}}, \widetilde{w}, \omega \right), \end{aligned} \tag{32}$$

where m^x_{ia} selects atoms with finite target and predicted coordinates, $\mathcal{I}_x = \{i : \sum_a m^x_{ia} > 0\}$, and the superscript ctr denotes centering under this mask. For either coordinate array, centering subtracts $\sum_a m^x_{ia} X_{ia} / \sum_a m^x_{ia}$ from each retained atom. This auxiliary loss removes translation but retains dependence on orientation in the supplied coordinate frame. It is distinct from the rigidly aligned RMSD used for evaluation. The distance-based losses depend only on atom-pair distances. The near-pair mask is

$$m^{\mathrm{near}}_{iab} = m^{TS}_{iab} \mathbf{1}[D^{TS}_{iab} \le 8\,\text{Å}].$$

The residual loss $\mathcal{L}_{\mathrm{res}}$ implements the squared residual reconstruction term in Eq. (8), using the masks and pair weights above. Squares inside Avg are elementwise. The remaining distance, diffusion-noise, structural, bond, and path-regularization terms form $\mathcal{L}_{\mathrm{aux}}$. $\mathcal{L}_{\mathrm{bond}}$ is an unweighted six-class cross-entropy over bond actions from reactants to products, with endpoint-derived training labels.

$$\begin{aligned} \mathcal{L}_{\mathrm{aux}} = \mathcal{L}_{\mathrm{rec}} &+ 0.10\mathcal{L}_\epsilon + \mathcal{L}_{\mathrm{lddt}} + 0.05\mathcal{L}_q \\ &+ 0.002\mathcal{L}_{\mathrm{coord}} + 0.25\mathcal{L}_{\mathrm{near}} + \mathcal{L}_{\mathrm{bond}} + 0.05\mathcal{L}_{\mathrm{fix}} \\ &+ \mathcal{L}_{\mathrm{profile}} + \mathcal{L}_{\mathrm{valid}}. \end{aligned} \tag{33}$$

Joint objective. The complete transition-state objective is

$$\mathcal{L}_{TS} = \mathcal{L}_{\mathrm{res}} + \lambda_c \mathcal{L}_c + \mathcal{L}_{\mathrm{aux}}, \qquad \lambda_c = 0.10. \tag{34}$$

All transition-state components are optimized jointly from the first step. Path assignment supplies detached targets, while the corrected estimate backpropagates through both the denoiser and Corrector.

A.4 TRANSITION-STATE IMPLEMENTATION DETAILS

This subsection records the generator configuration, DDPM forward parameterization, 25-step DDIM inference procedure, and target-free candidate selection.

Model Configuration. The generator predicts $\mathbf{R}_0 = (\mathbf{D}^{TS} - \mathbf{D}^{R})/s_\Delta$. Equations (24), (33), and (34) specify the correction, auxiliary, and joint objectives. The configuration is summarized below; all three learning rate groups use the same warm-up.

Setting	Value
Atom / pair state dimensions	384 / 192
Denoising blocks / attention heads	2 / 8
Corrector hidden dimension	192
Diffusion schedule / timesteps	Cosine / 1,000
DDIM evaluations / candidates	25 per candidate / 4
Residual / distance / coordinate scales	$s_\Delta = 2$ / $s_D = 10$ / $s_x = 2$ Å
Within-reactant distance basis	32 Gaussian centers over $[0, 40]$ Å
Denoiser learning rate	5×10^{-6}
Corrector learning rate	2×10^{-5}
Correction-scale head learning rate	2.5×10^{-6}
Learning rate warm-up	100 steps

Standard DDPM Kernel and Residual Coordinates. The forward process applies the fixed-variance DDPM kernel to the reactant-relative distance residual $\mathbf{R}_0$:

$$\begin{aligned} q(\mathbf{R}_t \mid \mathbf{R}_{t-1}) &= \mathcal{N}(\sqrt{\alpha_t}\mathbf{R}_{t-1}, \beta_t I)\,, \\ q(\mathbf{R}_t \mid \mathbf{R}_0) &= \mathcal{N}\big(\sqrt{\bar{\alpha}_t}\mathbf{R}_0, (1-\bar{\alpha}_t)I\big)\,. \end{aligned} \tag{35}$$

The Gaussian is defined on the independent valid upper-triangular entries. Each sampled noise value is mirrored to the corresponding lower-triangular entry; diagonal and invalid entries are zero. Thus I in Eq. (35) is the identity on the independent pair coordinates, not on all entries of the symmetric matrix. Enzyme-reactant interaction memory enters every denoising block, and the learned reaction path correction is applied through Eq. (5). The GFN2-xTB path in Eqs. (2) and (13) supplies training targets for these components.

Pairwise reconstruction places greater weight on short-range pairs, reactant bonds, and neighborhoods containing endpoint bond or charge changes. These endpoint-derived quantities affect only transition-state training losses and are not available to downstream predictors. Training also uses self-conditioning to couple successive denoising states. With probability 0.35, we first sample the implementation index $\tau' = \min(\tau + \delta, T-1)$, where $\delta = \mathrm{round}((T-1)/(25-1))$, and set $t' = \tau' + 1$. When $t' > t$, a prediction at t' is advanced deterministically to t; the resulting state replaces $\mathbf{R}_t$, its implied noise becomes the noise target, and the detached clean prediction provides self-conditioning. If this branch is not selected, with conditional probability 0.5 we use an estimate from the same state with gradients detached; otherwise, the self-conditioning input is zero.

Deterministic Sampling. At inference, the transition-state generator produces four distance matrix candidates through 25-step deterministic DDIM sampling. We use $T = 1000$ diffusion steps and the zero-based implementation grid $\tau_s = \mathrm{round}((T-1)(S-s)/(S-1))$, with $S = 25$. Let $t_s = \tau_s + 1$ denote the corresponding diffusion time in Eq. (4); thus $t_1 = T > t_2 > \cdots > t_S = 1$, and $\mathbf{R}_0$ always denotes the clean target. The model and schedule arrays use implementation index $t_s - 1$. Let $\mathcal{J} = \{1, \ldots, J\}$, with $J = 4$, index the candidates. For each $j \in \mathcal{J}$, sampling starts from a symmetric Gaussian residual $\mathbf{R}_{t_1,j}$. DDIM updates are deterministic conditional on this draw; distinct draws provide four candidates.

At sampling step s, let $\mathcal{M} = \mathcal{M}(\mathcal{E}, \mathcal{R})$. The predictor and Corrector in Eq. (5) give

$$\begin{aligned} \widetilde{\mathbf{R}}_{0,j}^{(t_s)} &= f_\phi(\mathbf{R}_{t_s,j}, t_s; \mathcal{M}, \mathbf{R}_{\mathrm{sc},s,j})\,, \\ (\mathbf{C}^r_{\phi,s,j}, \mathbf{C}^a_{\phi,s,j}) &= g_\phi\Big(\widetilde{\mathbf{R}}_{0,j}^{(t_s)}, t_s; \mathcal{M}\Big)\,, \\ \widehat{\mathbf{R}}_{0,j}^{(t_s)} &= \widetilde{\mathbf{R}}_{0,j}^{(t_s)} + \mathbf{C}^r_{\phi,s,j} + \mathbf{C}^a_{\phi,s,j}. \end{aligned} \tag{36}$$

Here we expose the optional self-conditioning input, suppressed in the main text: $\mathbf{R}_{\mathrm{sc},1,j} = 0$ at the first step and $\mathbf{R}_{\mathrm{sc},s,j} = \widehat{\mathbf{R}}_{0,j}^{(t_{s-1})}$ for $s > 1$. The noise implied by the corrected clean prediction is

$$\widehat{\boldsymbol{\epsilon}}_{t_s,j} = \frac{\mathbf{R}_{t_s,j} - \sqrt{\bar{\alpha}_{t_s}}\widehat{\mathbf{R}}_{0,j}^{(t_s)}}{\sqrt{1-\bar{\alpha}_{t_s}}}. \tag{37}$$

Using the next scheduled time gives the update in Eq. (6):

$$\mathbf{R}_{t_{s+1},j} = \sqrt{\bar{\alpha}_{t_{s+1}}}\widehat{\mathbf{R}}_{0,j}^{(t_s)} + \sqrt{1-\bar{\alpha}_{t_{s+1}}}\widehat{\boldsymbol{\epsilon}}_{t_s,j}, \tag{38}$$

for $s = 1, \ldots, S-1$. The main text abbreviates this next scheduled time as $t-1$. After the final evaluation, set $\widehat{\mathbf{R}}_{0,j} = \widehat{\mathbf{R}}_{0,j}^{(t_S)}$. The predicted transition-state distances are recovered in physical units, with the implementation's elementwise distance clipping:

$$\widehat{\mathbf{D}}_j^{TS} = \mathrm{clip}\Big(\mathbf{D}^R + s_\Delta \widehat{\mathbf{R}}_{0,j}, 0, 40\,\text{Å}\Big). \tag{39}$$

All residual matrices are symmetrized, set to zero on the diagonal, and masked to valid atom pairs throughout sampling.

Target-Free Selection. The four samples provide alternative distance matrix hypotheses for product free transition-state prediction. We select one hypothesis using a target-free score that combines three-dimensional reconstruction discrepancy, local reactant-distance preservation, and reactant-bond consistency:

$$\begin{aligned} j^\star &= \arg\min_{j\in\mathcal{J}} c_{\mathrm{str}}\left(\widehat{D}_j^{TS}; D^R, \mathcal{R}\right), \\ \widehat{D}^{TS} &= \widehat{D}_{j^\star}^{TS}. \end{aligned} \tag{40}$$

This score is deterministic and has no learned parameters. Let $D^{\mathrm{MDS}}(\widehat{D})$ denote the pairwise distances of a three-dimensional reconstruction of $\widehat{D}$ obtained by classical multidimensional scaling (MDS), $\Omega = \{(a,b) : a \neq b\}$, $\Omega_R = \{(a,b) \in \Omega : D_{ab}^R \leq 2.4\,\text{Å}\}$, and $\mathcal{B}_R$ be the reactant bond set. We use

$$\begin{aligned} c_{\mathrm{str}}(\widehat{D}; D^R, \mathcal{R}) = \frac{1}{|\Omega|}\sum_{(a,b)\in\Omega}\left|\widehat{D}_{ab} - D_{ab}^{\mathrm{MDS}}(\widehat{D})\right| \\ + \frac{0.25}{|\Omega_R|}\sum_{(a,b)\in\Omega_R}\left|\widehat{D}_{ab} - D_{ab}^R\right| \\ + \frac{0.5}{|\mathcal{B}_R|}\sum_{(a,b)\in\mathcal{B}_R}\left[\widehat{D}_{ab} - 2.4\,\text{Å}\right]_+, \end{aligned} \tag{41}$$

where a term is omitted when its index set is empty. The selected distance matrix is cached as the transition-state pair condition for product generation. Selection uses the reactant-side inference inputs.

A.5 Downstream Task Details

Both downstream tasks use the same trained transition-state generator. Product generation encodes the distance matrix selected from four Gaussian-initialized samples over SMILES token pairs, whereas compatibility ranking uses one denoising and correction evaluation at the zero-residual reactant anchor. We first describe the two interfaces, then present the ablation results.

Product Generation Interface

Transition-State Encoding. Let $A_{ia} \in \{0,1\}$ indicate that reactant token i is aligned to atom a. The normalized pooling weight is

$$w_{ia} = \frac{A_{ia}}{\max(1, \sum_b A_{ib})}. \tag{42}$$

Thus, a token aligned to multiple atoms averages their contributions, while special tokens and separators have zero atom weight. The alignment is determined entirely by reactant tokenization. For

each atom pair, the selected transition-state distance and its reactant-relative change are expanded with separate 32-dimensional Gaussian bases. These bases implement the Gaussian expansions described in Section 3.2. The normalized matrix $\mathbf{W} = [w_{ia}]$ maps each feature channel to token pairs by $\mathbf{W}(\cdot)\mathbf{W}^\top$, as in Eq. (9). The separate $\mathrm{MLP}_{\mathrm{abs}}$ and $\mathrm{MLP}_{\mathrm{rel}}$ projections are then summed to produce $\mathbf{B}^{TS}$. The resulting pairwise bias is clipped and added to self-attention over reactant tokens. Unaligned token pairs receive zero bias.

Product Architecture, Training, and Controls. The product model uses a SMILES encoder with 384 hidden dimensions, 4 layers, and 6 attention heads, and a Transformer decoder with 4 layers and 6 attention heads. USPTO-MIT pretraining maps mixed reactants and reagents to the product, after which the SMILES encoder and decoder initialize EnzymeMap fine-tuning. USPTO-MIT pretraining uses reaction records alone; enzyme and transition-state inputs are introduced during EnzymeMap fine-tuning. Both stages use the Molecular Transformer tokenizer, one fixed vocabulary, and one special-token mapping.

The fine-tuning encoder appends at most 32 residues from the same enzyme. Active- or binding-site annotations center this set, and nearby pocket residues preserve local context. Each example uses residues from one enzyme. Transition-state distances and reactant-relative changes use separate sets of 32 fixed Gaussian bases and projection heads. Each projected bias has scale 1.0, initial coefficient 0.5, and clipping bound 4.0. The decoder cross-attends to the resulting encoder states and generates the product. Beam search uses width 5, a maximum length of 224, and no length penalty; Top-k evaluation retains beam order, including repeated canonical products.

Matched transfer controls retain the same input width and trainable layers. ZymPro\E removes site tokens while retaining the enzyme-conditioned transition state. ZymPro\G (G: Geometry) removes only the TS geometry bias, retaining enzyme/site tokens, the encoder–decoder, pretraining, and the training protocol. ZymPro\T uses a product-endpoint pair bias. ZymPro\P removes USPTO-MIT pretraining of the product encoder and decoder.

The product model uses the following stage-specific settings.

Setting	USPTO-MIT pretraining	EnzymeMap fine-tuning
Learning rate	2×10^{-4}	10^{-4}
Dropout	0.10	0.15
EOS weight	2.0	1.0
Enzyme / TS conditioning	Absent	Present

Unless stated otherwise, neural models use AdamW with $\beta_1 = 0.9$, $\beta_2 = 0.95$, and weight decay 10^{-3}. Transition-state generation and product generation use batch size 24. Transition-state training uses validation early stopping with patience 10, and product training uses patience 16. Ranking-specific settings are given in Appendix A.7. The minimum-step gates and validation metrics are retained in the corresponding training configurations.

Endpoint-Supervised Control. ZymPro\T retains the enzyme-reactant encoder, interaction memory, and diffusion backbone, learning product-endpoint residuals with the same atom ordering, masks, and training set normalization:

$$\mathbf{R}_0^P = \frac{\mathbf{D}^P - \mathbf{D}^R}{s_\Delta}, \qquad \widehat{D}^P \sim p_{\phi_T}(D^P \mid \mathcal{E}, \mathcal{R}). \tag{43}$$

TS targets, path correction supervision, and path-energy losses are removed. All modules previously trained with these signals are retrained from the common generic initialization, matching reaction records, splits, training budgets, seeds, and validation selection. Inference uses only the enzyme and reactants. The same RBF projection maps $\widehat{D}^P$ to B^P for the USPTO-pretrained product decoder. For compatibility, endpoint atom/pair states from the same reactant anchor replace TS evidence in the unchanged transition-state readout, with correction features set to zero. Chemical validity and unchanged-reactant rates are separate generation diagnostics.

Compatibility Ranking Interface

Direct Enzyme–Reactant Scoring. The direct scorer uses fixed reactant global, pooled atom/pair, and metadata features and enzyme global, site, context, and structural features. Three projected

global channels compare semantic, site, and structural representations. After one residue-attention update with a bounded pair bias, softmax matching compares atoms with site residues, components with context and full-sequence tokens, and the global reactant with full-sequence tokens. Interaction MLPs pool 0.75 times the valid-token mean plus 0.25 times the maximum. Positive learned weights combine channels; local and full-sequence contributions bounded by $0.5\tanh(\cdot)$ are added to the global score to form u_c^{dir}. The token limits are 32 site residues, 16 context tokens, and 32 full-sequence tokens per enzyme.

Transition-State Evidence Extraction. For candidate enzyme $\mathcal{E}_c$, let $\mathcal{M}_c = \mathcal{M}(\mathcal{E}_c, \mathcal{R})$ be the interaction memory. The frozen transition-state feature extractor evaluates the denoiser once at the reactant anchor $\mathbf{R} = \mathbf{0}$, using implementation index $\tau = 0$ ($t = 1$ in the forward-process notation), followed by one path correction. It returns the initial and corrected residual estimates, $\widetilde{\mathbf{R}}_{0,c}$ and $\widehat{\mathbf{R}}_{0,c}$, and the denoiser's atom and pair states. Their changes are $\delta H_c = H_{R,c}^{\mathrm{out}} - H_{R,c}$ and $\delta B_c = B_{T,c}^{\mathrm{out}} - B_{T,c}^{\mathrm{in}}$, where $B_{T,c}^{\mathrm{in}}$ is the denoiser's initial pair state. The applied correction is $\Delta_c = \widehat{\mathbf{R}}_{0,c} - \widetilde{\mathbf{R}}_{0,c}$. Let Mean_A average over valid atoms and Mean_P average over valid off-diagonal atom pairs. For a scalar pair matrix X, let $\rho(X)$ concatenate its signed mean, absolute mean, and root-mean-square value over the same valid pairs. The transition-state evidence is

$$\begin{aligned}\mathbf{e}_c^{TS} = \mathrm{CONCAT}\big(&\mathrm{Mean}_A(\delta H_c), \mathrm{Mean}_P(\delta B_c),\\ &\rho(\widetilde{\mathbf{R}}_{0,c}), \rho(\widehat{\mathbf{R}}_{0,c}), \rho(\Delta_c)\big).\end{aligned} \tag{44}$$

A separate MLP maps $\mathbf{e}_c^{TS}$ to a scalar u_c^{TS}. It consists of layer normalization, a 192-unit hidden layer, SiLU, dropout 0.1, and a scalar output layer. The output layer starts at zero.

Compatibility Scoring and Training. The compatibility logit is the sum of the two branch scores:

$$u_c = u_c^{\mathrm{dir}} + u_c^{TS}. \tag{45}$$

Candidate enzymes are ranked by decreasing u_c. Each training group g contains one observed enzyme and three sampled alternatives. With $u_{g,+}$ denoting the observed enzyme's logit, the ranking objective is the mean group cross-entropy at temperature one:

$$\mathcal{L}_{\mathrm{rank}} = -\frac{1}{|\mathcal{G}|}\sum_{g\in\mathcal{G}} \log \frac{\exp(u_{g,+})}{\sum_{j\in g}\exp(u_{g,j})}. \tag{46}$$

The pretrained feature extractors and the interaction encoder, denoiser, and path corrector in the transition-state branch remain frozen. The direct scorer's projections, attention, channel weights, and scoring MLPs are trained from scratch jointly with the transition-state readout. Group membership enters the loss, while each branch computes its candidate score independently.

For evidence-decomposition controls, ZymPro\R uses u_c^{dir} alone and ZymPro\D uses u_c^{TS} alone. The full model adds both scores. ZymPro\C removes the Corrector while retaining the interaction and denoising paths; uncorrected denoising outputs provide product conditioning and ranking evidence, with correction-derived features set to zero. Candidate groups and the ranking objective are shared across these controls.

Ablation Results and Bond Rearrangement Metrics

Tables 6, 7, and 8 report the results underlying Fig. 7.

Table 6: Ablation study of enzyme compatibility ranking. All metrics except EF are percentages.

Variant	Top-1	MRR	Hits@2	ROC-AUC	PR-AUC	B_{85}	B_{20}	$EF_{0.05}$	$EF_{0.1}$
ZymPro\T	$49.28_{2.12}$	$65.17_{2.06}$	$66.34_{2.87}$	$65.42_{2.93}$	$46.16_{2.84}$	$62.37_{3.91}$	$60.24_{2.88}$	$2.32_{0.21}$	$2.21_{0.12}$
ZymPro\R	$57.14_{1.32}$	$74.63_{0.76}$	$79.18_{1.08}$	$76.39_{0.85}$	$54.16_{1.62}$	$79.41_{2.16}$	$68.59_{1.79}$	$2.84_{0.12}$	$2.73_{0.08}$
ZymPro\D	$55.68_{1.39}$	$73.91_{0.82}$	$78.16_{1.12}$	$75.53_{0.94}$	$52.86_{1.73}$	$78.34_{2.11}$	$67.23_{1.72}$	$2.77_{0.13}$	$2.65_{0.09}$
ZymPro\C	$\underline{59.27}_{1.29}$	$\underline{76.14}_{0.75}$	$\underline{81.36}_{1.06}$	$\underline{77.82}_{0.93}$	$\underline{56.67}_{1.61}$	$\underline{81.53}_{2.04}$	$\underline{70.48}_{1.66}$	$\underline{2.96}_{0.10}$	$\underline{2.83}_{0.08}$
ZymPro	$65.11_{1.25}$	$79.80_{0.79}$	$85.16_{0.96}$	$81.58_{0.90}$	$62.88_{1.76}$	$86.60_{2.07}$	$76.91_{1.69}$	$3.25_{0.09}$	$3.09_{0.07}$

Table 7: Ablation study of product generation. Accuracies are percentages.

Variant	Top-1 ↑	Top-3 ↑	Top-5 ↑
ZymPro\T	$40.28_{2.14}$	$53.76_{2.08}$	$61.31_{2.92}$
ZymPro\E	$44.17_{1.08}$	$59.34_{1.05}$	$64.72_{0.98}$
ZymPro\P	$22.36_{0.93}$	$37.18_{1.06}$	$44.27_{1.02}$
ZymPro\C	$47.08_{1.12}$	$63.23_{1.04}$	$68.54_{1.01}$
ZymPro\G	$\underline{48.55}_{1.03}$	$\underline{63.81}_{1.06}$	$\underline{68.78}_{1.01}$
ZymPro	$53.17_{1.04}$	$67.22_{0.98}$	$72.44_{0.94}$

Bond Rearrangement Directional Consistency. We construct a BRDC evaluation subset of 658 reactions from the rule-disjoint TS test set, each containing both bond formation and bond breaking. Let $\mathcal{F}_i$ and $\mathcal{B}_i$ denote the atom-mapped heavy-atom bond pairs that form and break. With the same subset $\mathcal{I}_{\mathrm{BRDC}}$ and $N_{\mathrm{BRDC}} = 658$ for every variant, we compute

$$F_i = \frac{1}{|\mathcal{F}_i|} \sum_{(a,b)\in\mathcal{F}_i} \mathbf{1}[\widehat{D}^{TS}_{i,ab} < D^{R}_{i,ab}],$$

$$B_i = \frac{1}{|\mathcal{B}_i|} \sum_{(a,b)\in\mathcal{B}_i} \mathbf{1}[\widehat{D}^{TS}_{i,ab} > D^{R}_{i,ab}],$$

$$\text{Formation} = \frac{100}{N_{\mathrm{BRDC}}} \sum_{i\in\mathcal{I}_{\mathrm{BRDC}}} F_i, \qquad \text{Breaking} = \frac{100}{N_{\mathrm{BRDC}}} \sum_{i\in\mathcal{I}_{\mathrm{BRDC}}} B_i,$$

$$\text{BRDC} = \frac{100}{N_{\mathrm{BRDC}}} \sum_{i\in\mathcal{I}_{\mathrm{BRDC}}} \frac{F_i + B_i}{2}.$$

Formation measures shortening in the forming pairs; Breaking measures lengthening in the breaking pairs. The inequalities are strict, so an unchanged distance does not count as a correct direction. Each reaction has equal weight, and formation and breaking have equal weight within a reaction, regardless of their bond counts. The metrics assess the direction, rather than the magnitude, of the distance change.

For bond localization, forming-bond AUROC ranks initially nonbonded heavy-atom pairs by relative shortening, $(D^R - \widehat{D}^{TS})/D^R$, with newly formed bonds as positives. Breaking-bond AUROC ranks initially bonded pairs by relative lengthening, $(\widehat{D}^{TS} - D^R)/D^R$, with broken bonds as positives. We average the two AUROCs within each reaction and then across the same 658 reactions, displaying the result multiplied by 100.

Table 8: Ablation of bond rearrangement directions and localization. AUROC is multiplied by 100; other values are percentages.

Variant	BRDC ↑	AUROC ↑	Formation ↑	Breaking ↑
ZymPro	84.0 ± 1.5	79.81 ± 1.14	94.7 ± 1.2	73.4 ± 2.8
ZymPro\A	80.5 ± 1.7	77.42 ± 1.26	92.5 ± 1.4	68.5 ± 3.1
ZymPro\B	78.0 ± 1.9	75.19 ± 1.38	90.0 ± 1.6	66.0 ± 3.4
ZymPro\C	75.0 ± 2.1	71.24 ± 1.48	87.5 ± 1.9	62.5 ± 3.8

A.6 Leakage Control and Evaluation Protocol

The data partitions are established before model training. Transition-state generation and compatibility ranking use rule-disjoint evaluation, while product generation uses product-disjoint evaluation. The two downstream tasks share one trained transition-state generator. Reaction-language pretraining is filtered against the protected EnzymeMap products.

Source Records. The processed EnzymeMap source contains 349,458 rows, 214,830 unique reaction indices, and 491 valid reaction rule identifiers. After enzyme assignment and product validation, the base record pool contains 142,608 records. Each task applies its own eligibility requirements and split assignment, so task datasets can share records and their sizes are not additive.

Rule-Disjoint Evaluation. Transition-state generation and enzyme compatibility ranking follow the same rule-disjoint assignment, with shared records assigned to the same split. The compatibility dataset includes proteins of at most 1,500 residues with complete enzyme and molecular features. Enzymes and reactants are encoded separately with $D_{RE} = 0$; candidate scoring does not use docked complexes. Training negatives may change across epochs, while validation and test candidate sets remain fixed. Each evaluation group fixes the reactants and contains the observed enzyme together with three alternatives sharing the first one, two, or three EC levels.

The ranking benchmark contains 87,849/6,464/5,255 training/validation/test groups spanning 178/84/75 rules.

Known-Positive Exclusion. For each canonical reactant-side input, we collect all enzymes annotated as positive across all source records and products. This complete known-positive set is excluded before EC-stratified sampling, including enzymes known to produce a different product from the same reactants. The exclusion applies to training-time resampling and the fixed validation and test candidate sets. Retained alternatives are unannotated candidates for that reactant-side input.

Top-1 and Hits@2 measure whether the observed enzyme ranks first or within the top two candidates; MRR averages its reciprocal rank across groups. ROC-AUC, PR-AUC, BEDROC, and EF are computed over all pooled enzyme-reaction pair scores. BEDROC uses $\alpha = 85, 20$; EF measures enrichment in the top 5% and 10%. All reported metrics except EF use a percentage scale.

Product-Disjoint Generation. Product generation uses exact canonical product groups after removing atom maps while retaining stereochemistry and all dot-separated components. Identical canonical reactant-enzyme inputs are assigned to the same split. Product components containing records used for transition-state training or validation are assigned to the product-training split. The transition-state generator is kept frozen when training either downstream model; the task-specific interfaces are described in Section A.5. An additional canonical-product check finds no overlap between the 28,974 transition-state training records and either the 10,313 product-validation records or the 9,519 product-test records.

Pretraining Isolation. We pool the official USPTO-MIT train, validation, and test records, remove reactions with more than 96 mapped precursor atoms or tokens outside the fixed vocabulary, and exclude exact canonical products assigned to EnzymeMap validation or test. The remaining records are repartitioned by connected components that close over identical canonical mixed sources and products. Consequently, records from each official partition may occur in each new partition, but canonical sources and products are disjoint across the new pretraining splits and protected EnzymeMap product overlap is zero. Enzyme and transition-state token-pair inputs are introduced during EnzymeMap fine-tuning.

Model Selection. We use five random seeds for ZymPro and all retrained baselines. Validation data determine model selection and inference budgets. Each run is evaluated once on the full test set after the corresponding task configuration is fixed.

A.7 Compatibility Baseline Adaptations

All compatibility models use the same enzyme-reactant candidate ranking protocol. We specify their common selection criterion and summarize the model-specific input representations and scoring mechanisms.

Common Protocol and Model Selection. All methods use the same rule-disjoint EnzymeMap partition and EC-matched candidate groups. Each of the 5,255 test groups contains one observed enzyme and one fixed alternative from each of the L1, L2, and L3 pools. These alternatives share the first one, two, and three EC levels with the observed enzyme, while differing at the next level. Training alternatives are resampled with deterministic epoch-dependent seeds, whereas validation and test alternatives remain fixed. All pools apply the known-positive exclusion in Section A.6. Model inputs are limited to reactants and the candidate enzyme; products, transition-state references, and EC labels are used for targets or benchmark construction, not as ranking inputs.

For all ranking methods, including ZymPro and the baselines, the reported checkpoint or tree count is selected by validation Top-1 accuracy. Task-specific scorers are fitted on the common rule-disjoint training split; frozen generic encoders that define a published input representation are retained.

The native reaction encoders of CLIPZyme (Mikhael et al., 2024) and ReactZyme (Hua et al., 2024) include the product side. To satisfy the common candidate-screening contract, we retain their released enzyme and reaction architectures, serialize only the reactant side, and retrain all task-specific components. The main table therefore uses their original names for reactant-only adaptations.

Model-Specific Adaptations. The following representations are used under the common protocol. Generic pretrained features remain frozen where indicated, while task-specific scorers are fitted on the common training partition.

Model	Representation and scoring mechanism
CLIPZyme	Released enzyme and reaction encoders, with reactant-only reaction serialization under the adaptation above.
FusionESP	Frozen ESM2 and MoLFormer features; released projection, normalization, cosine scoring, and cosine-MSE objective (Du et al., 2025).
EnzRank	Raw-sequence multi-kernel CNN and a 2,048-bit radius-2 chiral Morgan representation of the full reactant side (Upadhyay et al., 2023).
ESP	Frozen ESM-1b features and a 1,024-bit ECFP6 fingerprint; the released pretrained XGBoost scorer (Kroll et al., 2023) is fine-tuned on EnzymeMap, retaining 342 pretrained trees and 34 additional trees.
ProSmith	Released pooled gradient-boosting implementation: frozen ESM-1b features and mean-pooled ChemBERTa-77M-MTR logits, scored by XGBoost (Kroll et al., 2024).
ReactZyme	Released enzyme and reaction architectures, with reactant-only reaction serialization under the adaptation above.
EZSpecificity	Released no-structure setting: frozen ESM-2 residue features, a trainable D-MPNN with 3 layers, bidirectional cross-attention, and an MLP (Cui et al., 2025).

EZSpecificity's no-structure setting uses the same absence of a docked candidate complex as the other ranking inputs. Its enzyme residue embeddings are retained.

ZymPro Configuration. The ranking interface in Appendix A.5 uses width 96 and four attention heads for direct set matching, and width 192 for global projections and scoring hidden layers. AdamW uses learning rates 2×10^{-7} and 10^{-4} for the direct scorer and transition-state readout, respectively, with $(\beta_1, \beta_2) = (0.9, 0.999)$, weight decay 0.01, and gradient clipping at norm 1. Training uses batches of eight four-candidate groups for at most 30 epochs. Validation uses batches of 16 groups every 1,000 steps, with early stopping after eight evaluations without improvement. The selected checkpoint maximizes validation Top-1, with MRR and then Hits@2 breaking ties; its saved scoring parameters are used for evaluation.

Class-Wise Results. The test set contains 75 distinct reaction rule IDs. Grouping reactions by their complete rule-ID combination yields 77 signatures: 75 single-rule and 2 multi-rule signatures. Figure 4 includes the 21 classes with at least 20 reactions, ordered by class size and arranged in three columns. Each class compares ZymPro with the baseline attaining the highest Top-1 score; the remaining baselines appear as gray points. The complete test set results are reported in Table 2.

A.8 Product Baseline Adaptations

This subsection specifies the matched product generation contract before describing model-specific adaptations. Every baseline shares the protected pretraining corpus, EnzymeMap split, exact product criterion, and validation-only checkpoint selection; only its defining generation mechanism is retained.

Common Protocol. All product baselines are initialized randomly, pretrained on the same leakage-controlled USPTO-MIT split, and fine-tuned on the same EnzymeMap training split. We retain each method's defining encoder and generation mechanism, with the configurations below. Checkpoints are selected only on the corresponding validation split. Canonical product exact match is computed over all 9,519 test examples; samples that a graph generator cannot represent remain in the denominator.

Model-Specific Adaptations. The configurations below specify the adaptations to the common product task. P/F denotes USPTO-MIT pretraining/EnzymeMap fine-tuning; E, D, d, f denote en-

coder depth, decoder depth, hidden width, and feed-forward width. All models use eight attention heads. Batch budgets are per GPU, and "tok." denotes token-budget batching.

Model	Architecture and dropout	Optimization and batch budget (P/F)
Molecular Transformer (Schwaller et al., 2019)	$E = D = 4$, $d = 256$, $f = 2048$; dropout 0.1	Adam; Noam$(2, 8000)/(0.5, 400)$; 2,048 tok., accumulation 3
EC3-MT	Matched Molecular Transformer; level-3 EC token in fine-tuning	Matched Molecular Transformer
Chemformer (Irwin et al., 2022)	BART; $E = D = 6$, $d = 512$, $f = 2048$; dropout 0.1	AdamW; Noam$(1, 8000)/(0.2, 1000)$; batch 512/384
ReactionT5 (Sagawa & Kojima, 2025)	$E = D = 6$, $d = 512$, $f = 1408$; gated GELU; dropout 0.1	AdamW; Noam$(1, 8000)/(0.25, 2000)$; batch 512/384
Graph2SMILES (Tu & Coley, 2022)	4 graph layers, $E = D = 6$, $d = 256$, $f = 2048$; dropout 0.3	AdamW; Noam$(4, 8000)/(1, 500)$; 8,192/32,768 tok., accumulation 2/1
NERF (Bi et al., 2021)	Electron redistribution; depth 6, $d = 256$; dropout 0.1; VAE weight 0.1	AdamW; LR $5\times10^{-4}/10^{-4}$; warmup 320/32 steps; batch 1,024
ReactionSink (Meng et al., 2023)	Depth 6, $d = 256$; dropout 0.1; 4 Sinkhorn heads, 3 iterations; VAE weight 0.1	AdamW; LR $2 \times 10^{-4}/10^{-4}$; warmup 1,000/500 steps; batch 256

Noam(c, w) uses learning rate $\eta_t = c\, d^{-1/2} \min(t^{-1/2}, t\, w^{-3/2})$, with warmup w measured in optimizer steps. Adam uses $(\beta_1, \beta_2) = (0.9, 0.998)$ and zero weight decay; AdamW uses $(0.9, 0.95)$ and weight decay 10^{-3}. Gradient accumulation is one unless specified. NERF and ReactionSink use linear warmup followed by a constant learning rate. The five autoregressive models use beam width 5, maximum length 224, and zero length penalty. NERF and ReactionSink draw five VAE samples at temperatures $0.7(1.3)^j$, $j = 0, \ldots, 4$, retaining generation order.

Following the original NERF code, both baselines use a fixed source-atom graph and a decoder without stereochemical prediction, whereas evaluation requires all product components and stereochemistry. Unrepresentable samples count as failures. These constraints may contribute to their low exact-match accuracy on the product-disjoint EnzymeMap split. Their original USPTO-MIT scores use a different evaluation setting, including main-product matching (Bi et al., 2021; Meng et al., 2023).

Class-Wise Results. The product analysis groups 34 EC level-2 classes. Figure 5 retains the 18 classes with at least 20 reactions and displays one raincloud per method in separate Top-1, Top-3, and Top-5 panels. Points represent equally weighted EC classes; half violins show the distributions of class-level accuracies, and boxes show the median and interquartile range. Whiskers extend to the most extreme observed scores within 1.5 interquartile ranges of the box. Each method uses the same color across all three panels. Scores are aggregated from the retained predictions. Aggregate performance is reported in Table 3.

A.9 Transition-State Generation Evaluation

We evaluate geometric agreement with the computed path references under product free inference, comparing the full enzyme-conditioned generator with reactant-only baselines and an enzyme-free control.

Baselines and Enzyme-Free Control. OA-ReactDiff conditions on reactant and product structures (Duan et al., 2023) and is therefore outside this product free comparison. TSDiff (Kim et al., 2024) and GoFlow (Galustian et al., 2025) encode two-sided reaction graphs in their original formulations. We retrain both with reactant-only inputs and retain their respective generation mechanisms. Their names in Table 9 and Figures 6 and 8 refer to these adaptations. We additionally evaluate ZymPro\N (N: no enzyme), which removes both model enzyme conditioning and the enzyme pocket from the training energy calculation, evaluating their joint contribution. This differs from ZymPro\E, which removes site tokens only from the downstream product model.

Evaluation Protocol. We use the frozen rule-disjoint test split. The full model and ZymPro\N share the same evaluation references from substrate–pocket cluster calculations. Each method generates

four candidates and applies the same target-free selector in Eq. (41), based on three-dimensional reconstruction consistency, local reactant-distance preservation, and reactant-bond consistency. Selection does not access products or reference transition states. Distance RMSE is computed over off-diagonal atom pairs for each reaction; coordinate RMSD uses optimal translation and rotation alignment. For ZymPro, coordinates are fitted to predicted distances, initialized from the reactant geometry. Errors are summarized as medians across reactions. Each method produces 3,142 predictions; one non-finite reference is excluded from target-dependent metrics, leaving 3,141 evaluable reactions.

Table 9: Transition-state generation on the rule-disjoint test split. Errors are medians in Å over 3,141 finite-reference reactions from 3,142 predictions. Only full ZymPro uses enzyme inputs; all other rows use reactants alone. NFE/cand. denotes model function evaluations per generated candidate.

Method	NFE/cand.	Dist. RMSE ↓	RMSD ↓
GoFlow (Galustian et al., 2025)	24	2.624	2.981
TSDiff (Kim et al., 2024)	5,000	4.635	5.209
ZymPro\N	25	2.214	2.403
ZymPro	25	1.763	1.938

Full achieves the lowest median errors, followed by ZymPro\N, GoFlow and TSDiff. We evaluate larger substrates across diverse enzymatic reactions using pocket-derived references without product inputs; the original TSDiff and GoFlow benchmarks use small gas-phase molecules with both reactant and product graphs (Kim et al., 2024; Galustian et al., 2025).

A.10 Additional Compatibility Analyses

We detail the EC-proximity candidate tiers and complement the ablation studies with denoising-state sensitivity analyses.

EC-Proximity Candidate Tiers. For a given reactant-side input, the fixed L1-L3 test candidate set contains alternatives at three EC-proximity levels. Table 10 lists these tiers and an additional EC-unmatched random control, L0, constructed separately from the test set. Pairwise comparisons use the same reactant input and award half credit to tied scores. The alternatives are benchmark candidates rather than experimentally verified non-catalytic enzymes.

Table 10: EC-proximity candidate tiers. L0 is an EC-unmatched control constructed outside the fixed L1-L3 test candidate set.

Candidate alternative	L0 random control	L1	L2	L3
Shared EC levels	0	1	2	3

Denoising-State Sensitivity.

The compatibility model computes transition-state evidence online because candidate sets change during training. The transition-state generator can be evaluated at different diffusion states, so we vary the diffusion time and residual supplied to the denoising operation before path correction. We label probes by the implementation index $\tau = t - 1$. Under the residual parameterization, $\tau = 0, \mathbf{R} = 0$ places the denoising input at the observed reactant-distance anchor, $D = D^R$, and uses the resulting transition-state evidence. The default state is fixed before test evaluation; the remaining configurations are controlled sensitivity diagnostics and do not participate in model selection. Table 11 compares the default state, direct-only scoring, and the alternative denoising states.

Table 11: Denoising-state sensitivity. All metrics except EF are percentages.

Probe	Top-1	MRR	Hits@2	ROC-AUC	PR-AUC	B_{85}	B_{20}	$EF_{0.05}$	$EF_{0.1}$
Direct only	57.14	74.63	79.18	76.39	54.16	79.41	68.59	2.84	2.73
$\tau = 0, \mathbf{R} = 0$	65.11	79.80	85.16	81.58	62.88	86.60	76.91	3.25	3.09
$\tau = 10, \mathbf{R} = 0$	60.83	77.12	82.53	79.31	58.34	83.19	73.02	3.02	2.89
$\tau = 100, \mathbf{R} = 0$	60.18	76.82	81.97	78.72	57.98	82.63	71.82	3.01	2.86
$\tau = 500, \mathbf{R} = 0$	58.92	75.88	80.62	77.43	55.87	80.58	69.73	2.93	2.78
$\tau = 999, \mathbf{R} = 0$	56.88	74.52	78.59	75.83	53.27	78.13	67.58	2.79	2.68
$\tau = 999, \mathbf{R} = \boldsymbol{\epsilon}$	57.81	75.12	79.52	76.57	54.53	79.02	68.48	2.87	2.73

A.11 Scope and Limitations

The transition-state references in this work are approximated by the highest-energy internal frames along substrate–pocket cluster paths computed with GFN2-xTB. Each record supplies one computed path, and geometric evaluation therefore measures agreement with these computational path-derived references rather than with experimentally characterized transition states.

Our primary goal is to study whether transition-state and reaction-path supervision can serve as useful process-level inductive bias for enzyme-conditioned prediction. Accordingly, ZymPro should not be interpreted as establishing experimentally exact transition states or complete enzyme-catalyzed reaction mechanisms. The path-derived transition-state estimate serves as an operational intermediate representation through which process-aware supervision is transferred to product generation and enzyme compatibility ranking. The downstream results are intended to evaluate the predictive utility of this representation, while the geometric analyses characterize its consistency with the computational supervision used in this work.

The model requires enzyme residue features, within-system structures, and site annotations, whose availability and quality affect its application. It does not require a docked enzyme–reactant complex. Future work may evaluate higher-level electronic-structure references, multiple competing reaction pathways, kinetic and energetic quantities, and experimental validation to further assess the physical fidelity of the generated transition-state estimates.

A.12 Reaction-to-Transition-State Trajectories

Figure 8 extends Figure 6 to ten reaction schemes and two-distance projections. All displayed model trajectories start from the same reactant geometry. Coral ZymPro curves show 25 sampled input states from the zero-residual reactant anchor and the final prediction. Dashed blue curves connect MolecularGSM frames to the path-derived TS reference. Baseline endpoints use the reactant-only adaptations in Appendix A.9, selecting one of four candidates by Eq. (41). Marker definitions follow Figure 6. Schemes omit water and protons and abbreviate cofactors.

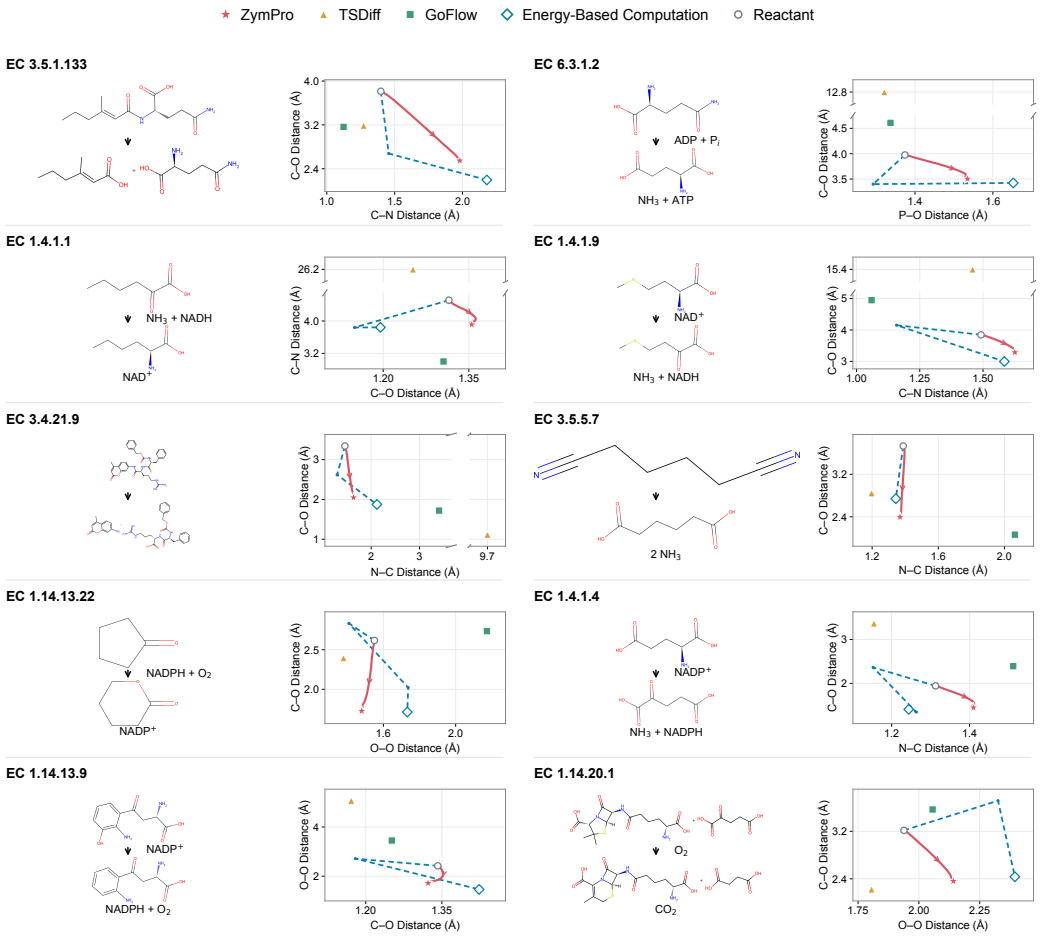

Figure 8: Reaction schemes, distance trajectories, and TS estimates from ZymPro, TSDiff, and GoFlow for ten cases. EC labels identify enzyme classes; markers and colors match Figure 6. Axis breaks indicate omitted distance ranges.