TACFORMER: TRANSFORMER-BASED GENERATIVE MODEL FOR TARGET-CONDITIONED *De Novo* PROTAC DESIGN

Anonymous authors
Paper under double-blind review

ABSTRACT

Targeted protein degradation (TPD) represents a fundamental shift in drug discovery, moving beyond traditional target inhibition toward the event-driven removal of disease-causing proteins. A prominent strategy for TPD is the use of proteolysis-targeting chimeras (PROTACs), heterobifunctional molecules that selectively route a protein of interest (POI) to the ubiquitin–proteasome system for degradation. While PROTAC design still relies heavily on trial-and-error synthesis and experimental testing, deep generative models offer a promising means to explore PROTAC chemical space *in silico* and generate novel molecules. However, many existing generative methods formulate PROTAC design primarily as a linker-generation problem, with POI-binding (warhead) and E3-ligase-binding components fixed in advance. This restricts exploration of the full PROTAC chemical space and often leaves the POI implicit through the predefined warhead, rather than using protein-level information as an explicit conditioning signal during generation. To address these limitations, we propose TACFormer: a transformer-based generative model for target-aware *de novo* PROTAC design. The model conditions a GPT-style molecular decoder on pretrained ESM-2 representations of POI sequences, allowing POI-level information to steer full-molecule generation. Our results show that TACFormer generates highly valid, novel, and chemically diverse PROTACs while showing consistent POI-dependent shifts in generated chemical space. In particular, molecules generated under specific target conditions preferentially populate chemical neighborhoods associated with POIs in the held-out dataset, while extrapolating to novel chemotypes. By natively integrating biological context into the generative process, TACFormer moves PROTAC design toward a more biology-informed exploration of degrader chemical space.

1 INTRODUCTION

Targeted protein degradation (TPD) has emerged as a promising paradigm in drug discovery, replacing sustained occupancy of a functional site with the event-driven removal of the target protein itself (Békés et al., 2022). Because it does not require a druggable active site, TPD extends to targets long considered "undruggable". Proteolysis-targeting chimeras (PROTACs) are its most developed modality, with several having advanced into clinical trials (O'Brien Laramy et al., 2023): heterobifunctional molecules built from a warhead that binds a protein of interest (POI), a ligand that recruits an E3 ubiquitin ligase, and a linker joining the two. By binding both moieties simultaneously, a PROTAC brings them into proximity, inducing polyubiquitination of the POI, marking it for degradation by the proteasome.

Despite this promise, PROTAC discovery remains driven by iterative synthesis and experimental screening, and the search must be repeated for each new POI (Pike et al., 2026). Deep generative models offer an alternative by learning distributions over known chemistry and proposing molecules beyond existing libraries. Transformer-based chemical language models are well matched to the problem: they operate on simplified molecular-input line-entry systems (SMILES) (Weininger, 1988), whose non-injective syntax permits randomization as data augmentation (Arús-Pous et al., 2019), and they remain effective when only a few hundred examples are available (Moret et al.,

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.

2020). This latter point is particularly important as experimentally characterized degraders number only in the thousands.

Existing generative approaches for PROTAC design largely formulate the task as linker generation (Gharbi & Mercado, 2024), with the warhead and E3-ligase ligand fixed in advance. This confines the search to the linker, represents the POI only implicitly through the chosen warhead, and restricts *de novo* design to targets that already have known binders. The few methods that generate complete degraders introduce target information indirectly, through activity surrogates or docking-based rewards, rather than conditioning the generator on the target protein itself. To date, no method designs full PROTACs *de novo* from an explicit representation of the target.

In this work, we aim to address the gaps of existing works and design a method that (i) generates complete PROTAC molecules rather than restricting generation to the linker; (ii) integrates the biological representation of the POI as a conditioning signal; and (iii) operates effectively in low-data regimes inherent to PROTAC discovery. We summarize our contributions as follows:

- We introduce TACFormer, to our knowledge, the first generative model for *de novo*, full-molecule PROTAC design that conditions generation directly on a learned representation of the target protein. POI sequences are encoded with a pretrained ESM-2 model and supplied as an explicit conditioning signal to a decoder-only SMILES generator, so that target information steers generation of the full molecule.
- We address the low-data regime through two-stage training: we curate a corpus of 79,406 PROTAC-like molecules by filtering ZINC and ChEMBL for the physicochemical and bifunctional profile of degraders, pretrain a chemical prior on it, and fine-tune on reported PROTACs. We also contribute a normalization pipeline that harmonizes POI and E3-ligase annotations across the three source databases.
- We establish that POI carries a measurable chemical signature in PROTAC data, showing that POI-associated chemistry is globally heterogeneous but locally strongly organized.
- We propose an evaluation protocol for target-conditioned molecular generation that combines a likelihood-based measure (Δ_{NLL}) with nearest-neighbor purity in chemical space, and use it to characterize both where conditioning succeeds and where it fails, including a zero-shot evaluation on POIs unseen during training.

2 Related Work

Transformer-based molecular generation. Transformer architectures have been standard for *de novo* molecular design using string-based representations. Encoder–decoder models can handle sequence-to-sequence problems such as reaction prediction, retrosynthesis, and molecular optimization (Schwaller et al., 2019; Irwin et al., 2022). Encoder-only models learn bidirectional representations of molecular sequences and serve mainly predictive ends, including property and activity modeling. Decoder-only models use causal self-attention to factorize the sequence autoregressively, which maps directly onto generation and admits conditioning on physicochemical properties, scaffolds, or fragments. Conditioning on biological rather than chemical context is less common: (Chen & Bajorath, 2024) condition a SMILES decoder on protein-language-model embeddings to generate compounds with a target potency against a given protein, but consider single-ligand rather than bifunctional design. Across all three variants, pretraining on a broad molecular corpus and fine-tuning on a smaller domain-specific set is a robust route to transferring chemical representations into data-limited settings.

Generative PROTAC linker design. Most generative work on PROTACs treats the problem as linker design (Gharbi & Mercado, 2024). Early methods were general-purpose fragment-linking models trained on ZINC or ChEMBL and applied to degraders only as a demonstration, without PROTAC-specific training data (Imrie et al., 2020; Yang et al., 2020). Subsequent methods have been tailored to degraders through PROTAC-specific fine-tuning and objectives that reflect their distinct chemical space: string-based autoregressive generation steered by reinforcement learning (RL) (Guo et al., 2023), three-dimensional generation conditioned on the ternary-complex pose (Li et al., 2023; Neeser et al., 2023), and equivariant diffusion over atomic coordinates (Igashov et al., 2024). Each explores linker chemistry systematically but require the warhead and E3-ligase ligand

as input, confining the search to the linker subspace, and leaving the combinatorial space of the whole molecule unoptimized.

Full-molecule PROTAC generation and target information. A smaller body of work generates the complete degrader. (Nori et al., 2022) build PROTAC-like molecular graphs from an empty graph and introduce target specificity through RL, rewarding a surrogate model that predicts degradation activity for a given target. TRL-PROTAC (Dai & Zhu, 2024) generates POI and E3 ligands separately, joins them with a linker model, and optimizes the result against pharmacokinetic and degradation objectives, while (Shao et al., 2024) drive a language-model pipeline with paired structural and property constraints. SynPROTAC (Xu et al., 2026) assembles synthesizable degraders by reaction-based construction, but draws warheads and E3 ligands from the dataset, using target information through docking-based RL rather than a direct conditioning signal for the generator.

3 PRELIMINARIES

3.1 AUTOREGRESSIVE FORMULATION

PROTACs are represented as SMILES strings, allowing molecular generation to be formulated as autoregressive language modeling. For a molecule tokenized as $\mathbf{s} = (s_1, \ldots, s_T)$, the conditional distribution is factorized as $P_\theta(\mathbf{s} \mid p) = \prod_{t=1}^{T} P_\theta(s_t \mid s_{<t}, p)$, where θ denotes the model parameters and p the conditioning POI. The conditional model therefore learns a family of POI-dependent molecular distributions rather than a single unconditional distribution over degraders. Conditioning on p is implemented through its fixed embedding $\mathbf{z}_p$; accordingly, we use $P_\theta(\cdot \mid p)$ and $P_\theta(\cdot \mid \mathbf{z}_p)$ interchangeably. Tokenization and SMILES randomization are described in Appendix A.1.

3.2 PROTEIN REPRESENTATION USING ESM-2

Target information is represented using ESM-2 (Lin et al., 2023), a transformer protein language model that learns sequence-derived representations containing evolutionary and structural information. For each POI, we retrieve the corresponding amino-acid sequence from UniProtKB and encode it with the pretrained `esm2_t33_650M_UR50D` variant (33 layers, 650M parameters, hidden size 1280), yielding $\mathbf{H}_p = [\mathbf{h}_1, \ldots, \mathbf{h}_L]^\top \in \mathbb{R}^{L \times 1280}$ for a protein of length L after removing special tokens. We then mean-pool over residues to obtain a single fixed POI representation, $\mathbf{z}_p = L^{-1} \sum_{i=1}^{L} \mathbf{h}_i \in \mathbb{R}^{1280}$. Proteins longer than the 1022-residue processing window are encoded in overlapping windows and reconstructed to full length before pooling, ensuring that each residue contributes exactly once (Appendix A.3). ESM-2 remains frozen throughout.

3.3 CONDITIONAL DECODER

The molecular generator is a decoder-only transformer with eight pre-normalization blocks, a hidden dimension of $d_{\text{model}} = 512$, and eight attention heads (Appendix A.4); the same backbone is used for molecular pretraining, unconditional fine-tuning, and target-conditioned fine-tuning. The pooled protein representation is projected into the decoder hidden space as $\mathbf{c}_p = \mathbf{W}_c \, \mathrm{LN}(\mathbf{z}_p) + \mathbf{b}_c$, with $\mathbf{W}_c \in \mathbb{R}^{512 \times 1280}$.

We introduce a global gated condition: each block ℓ contains a trainable scalar γ_ℓ, and $\mathbf{c}_p$ is broadcast across token positions and added after the attention residual, such that $\mathbf{A}^{(\ell)} = \mathbf{X}^{(\ell)} + \mathrm{MHA}(\mathrm{LN}_1(\mathbf{X}^{(\ell)}))$ becomes

$$\widetilde{\mathbf{A}}^{(\ell)} = \mathbf{A}^{(\ell)} + \tanh(\gamma_\ell)\, \mathbf{c}_p, \qquad \mathbf{X}^{(\ell+1)} = \widetilde{\mathbf{A}}^{(\ell)} + \mathrm{MLP}\left(\mathrm{LN}_2\left(\widetilde{\mathbf{A}}^{(\ell)}\right)\right). \tag{1}$$

The same $\mathbf{c}_p$ is shared across all transformer blocks, while the trainable layer-specific gates γ_ℓ control the contribution of the protein condition at each depth. The conditional model is initialized from the pretrained molecular prior and augmented with the protein adapter and one gate per block. Since $\gamma_\ell = 0$ at initialization, $\tanh(\gamma_\ell)\mathbf{c}_p = \mathbf{0}$ and the model initially behaves identically to the unconditional decoder; conditioning is learned progressively during fine-tuning. Training uses teacher-forced causal language modeling on molecule–POI pairs.

4 EXPERIMENTS

4.1 DATA

PROTAC-like pretraining data. To learn a PROTAC-oriented molecular prior, compounds from ZINC20 (Irwin et al., 2020) and ChEMBL36 (Zdrazil et al., 2024) were filtered for physicochemical and structural properties characteristic of PROTAC-like molecules. After standardization, filtering, and deduplication, the final pretraining corpus contained 79,406 unique molecules. Detailed filtering criteria are provided in Appendix B.1.

PROTAC data. Two datasets were constructed by merging PROTAC-DB (Ge et al., 2025), TPDDB (Qin et al., 2026), and PROTAC-Pedia (London & Prilusky, 2020).

For unconditional generation, target annotations were discarded, and standardized SMILES were deduplicated globally, yielding 27,489 unique PROTACs. These were randomly split into training, validation, and test sets using an 80/10/10 ratio.

For conditional generation, the dataset was defined at the molecule–POI association level. POIs were normalized using UniProt accessions as protein identifiers and standardized gene names as target labels. The resulting dataset contained 30,053 associations across 408 POIs. WDR5 and DHFR were reserved exclusively for zero-shot evaluation, leaving 29,944 associations across 406 POIs for seen-target training and evaluation. The 20 most represented POIs, comprising 21,246 associations, formed the seen-POI evaluation panel, while the remaining 386 POIs were retained as training-only targets. Split construction and target statistics are provided in Appendix B.3.

4.2 EXPERIMENTAL SETUP

Experiments followed a shared-pretraining design. The decoder-only transformer was first pretrained on the PROTAC-like corpus to learn a molecular prior, which was then used to initialize two downstream models trained on PROTAC data.

The *unconditional model* was fine-tuned on the globally deduplicated PROTAC dataset without target information, learning a distribution $P_\theta(\mathbf{s})$ over complete PROTAC molecules. The *conditional model* was initialized from the same molecular prior and fine-tuned on molecule–POI associations using the protein-conditioning mechanism described in Section 3, learning $P_\theta(\mathbf{s} \mid \mathbf{z}_p)$.

Unconditional and conditional generation were compared using the same molecular backbone, initialization, and sampling budget. A total of 200,000 molecules were generated from each model. For the conditional model, samples were distributed evenly across the 20 seen evaluation POIs, yielding 10,000 molecules per POI.

The conditional model was evaluated under two target regimes. In the *seen-POI* setting, generation was conditioned on the 20 POIs represented during conditional training. In the *zero-shot* setting, the same selected checkpoint was used without further optimization and conditioned on WDR5 or DHFR, whose associated PROTAC chemistry was entirely excluded from conditional training. Training hyperparameters are provided in Appendix A, and generation settings in Appendix B.5.

4.3 EVALUATION

Let $\mathcal{G}$ denote the complete set of generated SMILES, $\mathcal{V} \subseteq \mathcal{G}$ the subset that can be successfully parsed and sanitized by RDKit, and $\mathcal{U}$ the set of distinct valid molecules obtained after canonicalization and deduplication of $\mathcal{V}$. For conditional generation, these sets are defined separately for each requested POI p and are denoted $\mathcal{G}_p$, $\mathcal{V}_p$, and $\mathcal{U}_p$, respectively. For the conditional model, metrics are additionally computed independently for each POI and summarized using an unweighted macro average across targets. Sampling and implementation details are provided in Appendix B.5.

Evaluation rationale. We first examine how POI identity is reflected in PROTAC chemical space using Morgan–Tanimoto similarity. Same- and different-POI relationships are evaluated both globally and within the CRBN and VHL E3-ligase backgrounds. Local POI organization is further assessed using nearest-neighbor purity, with a within-E3 label-permutation null as reference. This

analysis motivates evaluating conditional generation using both likelihood-based condition sensitivity and enrichment within local target-associated chemical neighborhoods, rather than similarity to a single target-specific chemical centroid. Details are provided in Appendix B.4, and the corresponding results are shown in Figure 3.

Generation quality. Validity, uniqueness, and novelty are defined as

$$\mathrm{V} = \frac{|\mathcal{V}|}{|\mathcal{G}|}, \qquad \mathrm{U} = \frac{|\mathcal{U}|}{|\mathcal{V}|}, \qquad \mathrm{Nov} = \frac{|\{g \in \mathcal{U} : g \notin \mathcal{R}_{\mathrm{train}}\}|}{|\mathcal{U}|}. \tag{2}$$

Validity measures the fraction of generated molecules that are valid, uniqueness measures the fraction of valid molecules that are distinct, and novelty measures the fraction of unique valid molecules that are absent from the training set.

For conditional generation, the same metrics are also computed separately for each POI p by replacing $(\mathcal{G}, \mathcal{V}, \mathcal{U})$ with the corresponding POI-specific sets $(\mathcal{G}_p, \mathcal{V}_p, \mathcal{U}_p)$. Conditional novelty is evaluated against the complete conditional training set, rather than only the training molecules associated with the requested POI, providing a stricter and more conservative measure of novelty.

We additionally report internal diversity to quantify the structural diversity among unique valid generated molecules. Given their 2048-bit Morgan fingerprints with radius 2, denoted by f_i, internal diversity is defined as

$$\mathrm{IntDiv} = 1 - \binom{|\mathcal{U}|}{2}^{-1} \sum_{i<j} T(f_i, f_j), \tag{3}$$

where $T(f_i, f_j)$ denotes the Tanimoto similarity between fingerprints f_i and f_j. Higher values therefore indicate greater structural diversity within the generated set.

Generated molecules are further characterized by molecular weight, TPSA, LogP, heavy-atom count, rotatable-bond count, aromatic and non-aromatic ring counts, hydrogen-bond acceptors, and hydrogen-bond donors. Fingerprint and property-computation details are provided in Details of the fingerprint representation and property calculations are provided in Appendix B.6.

POI-warhead analysis. Whole-molecule diversity does not necessarily reflect diversity at the POI-binding end of a PROTAC. We therefore decompose unique valid generated PROTACs into POI warhead, linker, and E3-ligand components using PROTAC-Splitter (Ribes et al.). For POI p, let $\mathcal{W}_p$ denote the processed warhead instances, $\mathcal{W}_p^*$ the corresponding unique canonical warheads, and $\mathcal{R}_p^W$ the set of unique warheads extracted from training PROTACs associated with the same POI. We report

$$\mathrm{WUniq}(p) = \frac{|\mathcal{W}_p^*|}{|\mathcal{W}_p|}, \qquad \mathrm{WNov}(p) = \frac{|\mathcal{W}_p^* \setminus \mathcal{R}_p^W|}{|\mathcal{W}_p^*|}, \tag{4}$$

together with internal diversity over the unique generated warheads,

$$\mathrm{WDiv}(p) = 1 - \binom{|\mathcal{W}_p^*|}{2}^{-1} \sum_{i<j} T(w_i, w_j). \tag{5}$$

PROTAC-Splitter processing and quality-control details are provided in Appendix B.9.

Conditional fidelity. Generation quality alone does not establish whether the decoder uses the protein condition. We therefore evaluate conditional fidelity using two complementary measures.

First, we evaluate real held-out molecules through a length-normalized negative log-likelihood,

$$\mathrm{NLL}(x \mid p) = -\frac{1}{T_x} \sum_{t=1}^{T_x} \log P_\theta \left(s_t \mid s_{<t}, \mathbf{z}_p\right). \tag{6}$$

For a held-out molecule x associated with POI p, let $\mathcal{P}(x)$ denote all of its valid target annotations and $\mathcal{T}$ the evaluated target panel. Conditional preference is measured by

$$\Delta_{\mathrm{NLL}}(x, p) = \frac{1}{|\mathcal{T} \setminus \mathcal{P}(x)|} \sum_{q \in \mathcal{T} \setminus \mathcal{P}(x)} \mathrm{NLL}(x \mid q) - \mathrm{NLL}(x \mid p). \tag{7}$$

A positive value indicates that the molecule receives higher likelihood under its correct POI condition than under competing conditions. All legitimate alternative targets of a multi-target molecule are excluded from the competing set.

Second, we evaluate whether generated molecules occupy held-out chemical neighborhoods associated with the requested target. Only valid, unique, training-novel generated molecules are used as queries. For query g generated under POI p, let $\mathcal{N}_k(g)$ denote its k nearest molecules in the held-out test reference pool under Morgan–Tanimoto similarity. Requested-POI purity is

$$\rho_k(g, p) = \frac{1}{k} \sum_{r \in \mathcal{N}_k(g)} \mathbf{1}\left[p \in \mathcal{P}(r)\right], \tag{8}$$

and is averaged over queries generated for the same POI. We report $k \in \{5, 10, 20\}$. The NLL and nearest-neighbor evaluations use held-out chemistry for the same target panel being evaluated; full reference-pool construction is described in Appendix B.8.

5 Results

5.1 Molecular pretraining and unconditional generation

Pretraining on the PROTAC-like data produces a molecular prior capable of generating predominantly valid and diverse large molecules. From 10,000 generated samples, the pretrained model achieves 98.18% validity, 91.99% uniqueness, and an internal diversity of 0.836.

We next compare the unconditional and conditional models using the same generation. Fine-tuning the molecular prior on PROTAC data without target information yields 98.78% validity, 50.20% uniqueness, and an internal diversity of 0.790. Introducing POI conditioning preserves similarly high validity while reducing uniqueness to 42.54% and internal diversity to 0.637.

5.2 Target-conditioned generation for seen POIs

We next evaluated whether conditioning on protein representations steers generation toward chemistry associated with the requested POI. Table 1 reports generation-quality and conditional-fidelity metrics for the 20 seen evaluation POIs.

Generation validity is consistently high across targets, with a macro average of 98.30% and a minimum of 96.26%. Novelty averages 91.59%, showing that most unique valid generated structures are absent from the conditional training chemistry. Uniqueness and internal diversity are more variable across conditions, with macro averages of 42.54% and 0.637, respectively.

Conditional fidelity is supported by both likelihood- and chemistry-based evaluation. Δ_{NLL} is positive for all 20 POIs, with a macro average of $+0.358$, indicating that held-out molecules consistently receive lower NLL under their associated POI condition than under competing conditions. The same target dependence is observed in generated chemical space: the mean 5-NN POI purity is 82.11%, decreasing to 78.37% at $k = 10$ and 72.41% at $k = 20$. Generated molecules are therefore preferentially located near held-out chemistry associated with the POI used for generation.

The strength of this effect varies substantially across targets. At $k = 5$, POI purity ranges from 23.85% for SMARCA4 to 98.93% for METTL14. High purity is not restricted to the most highly represented targets: for example, METTL14 reaches the highest 5-NN purity in the panel, whereas BRD4 reaches 64.22%. METTL14 also produces the most concentrated generated distribution, with the lowest uniqueness (16.77%) and internal diversity (0.394), while MAP4K1 has the highest internal diversity (0.722) and substantially lower 5-NN purity (57.63%). The relationship between conditional fidelity and generated diversity is examined further in Appendix D.

Component-level analysis further shows that conditional generation produces novel predicted POI-binding warheads. Across the 20 POIs, predicted POI-binding warheads have a mean target-warhead novelty of 89.18% relative to the corresponding training warheads. Figure 1 illustrates generated PROTACs for BRD4, MAP4K1, and BTK, the three POIs with the highest generated-warhead internal diversity, and shows multiple distinct predicted warhead scaffolds produced under each target condition.

Table 1: **Target-wise generation quality and conditional fidelity across the 20 seen evaluation POIs.** Generation-quality metrics are computed independently for each conditioning target. Δ_{NLL} measures the likelihood preference for the associated POI condition over competing conditions, while k-NN purity measures enrichment of the requested POI among the nearest held-out reference molecules. The macro row reports the unweighted mean across POIs and the weighted row aggregates according to the number of evaluated samples. Best per-target values in each column are shown in bold.

POI	Validity	Uniqueness	Novelty	Diversity	Δ_{NLL}	5-NN purity	10-NN purity	20-NN purity
AR	99.59%	49.48%	82.73%	0.689	0.246	97.18%	96.05%	95.12%
ESR1	99.28%	43.23%	84.37%	0.659	0.308	97.34%	96.49%	93.46%
CDK4	99.52%	33.54%	84.60%	0.580	**0.508**	96.86%	95.38%	93.71%
KRAS	99.16%	52.39%	88.32%	0.638	0.446	95.99%	95.02%	94.26%
EZH2	97.92%	43.64%	96.07%	0.644	0.269	76.26%	68.16%	51.56%
BTK	96.52%	49.99%	91.09%	0.716	0.342	88.82%	86.00%	81.08%
BRD4	96.33%	48.26%	93.01%	0.710	0.225	64.22%	57.60%	50.74%
BRAF	99.53%	32.65%	91.42%	0.589	0.297	91.81%	90.31%	86.07%
SMARCA4	98.14%	46.85%	92.65%	0.671	0.286	23.85%	23.13%	22.36%
MAP4K1	98.16%	47.24%	96.16%	**0.722**	0.345	57.63%	38.87%	22.95%
IRAK4	97.09%	**64.60%**	96.33%	0.709	0.433	91.46%	88.46%	82.23%
EGFR	97.11%	57.85%	91.72%	0.719	0.326	88.85%	82.65%	75.30%
BCL2L1	98.38%	41.99%	90.85%	0.674	0.348	85.41%	83.90%	78.52%
SMARCA2	97.91%	46.46%	92.46%	0.679	0.297	75.08%	72.91%	68.14%
METTL14	**99.83%**	16.77%	84.95%	0.394	0.484	**98.93%**	**98.78%**	**98.47%**
PTPN1	96.26%	28.47%	94.45%	0.603	0.461	78.47%	75.35%	73.86%
BCL2	97.87%	31.79%	95.24%	0.607	0.330	73.76%	70.79%	53.20%
SOS1	99.01%	43.64%	**96.44%**	0.663	0.350	77.17%	67.90%	52.07%
BCL6	98.92%	33.99%	93.63%	0.495	0.450	97.48%	97.42%	95.95%
MAPT	99.53%	37.88%	95.41%	0.590	0.412	85.62%	82.32%	79.15%
Macro	98.30%	42.54%	91.59%	0.637	0.358	82.11%	78.37%	72.41%
Weighted	98.30%	42.49%	91.78%	–	0.343	81.14%	77.04%	70.73%

5.3 Zero-shot generation for unseen POIs

Finally, we evaluated the same conditional model on WDR5 and DHFR. Generation quality remains high in terms of uniqueness and novelty, although validity differs between the two targets: 70.44% for WDR5 and 92.82% for DHFR (Table 2).

In contrast to the seen-target results, conditional fidelity does not transfer to either unseen POI. Δ_{NLL} is negative for both WDR5 (-0.053) and DHFR (-0.064), indicating that their held-out molecules do not receive a likelihood advantage under the corresponding unseen protein condition relative to competing conditions.

Nearest-neighbor purity likewise shows weaker target-specific enrichment than observed for seen POIs. Macro purity is 52.23% at $k = 5$, 53.50% at $k = 10$, and 56.08% at $k = 20$, rather than showing the strong local enrichment observed across the seen-target panel. Taken together, these results show that the conditional signal learned for targets represented during fine-tuning does not directly extend to WDR5 and DHFR under the present zero-shot setting.

Table 2: **Generation quality and conditional fidelity for the two unseen POIs.** WDR5 and DHFR were excluded from conditional model training and are supplied only through their pretrained protein representations at inference. Macro values are the unweighted means across the two targets.

POI	Validity	Uniqueness	Novelty	Scaffold novelty	Diversity	Δ_{NLL}	5-NN purity	10-NN purity	20-NN purity
WDR5	70.44%	**99.66%**	**99.97%**	**99.77%**	**0.739**	-0.053	**62.74%**	**59.78%**	55.83%
DHFR	**92.82%**	90.82%	98.79%	98.03%	0.679	-0.064	41.73%	47.23%	**56.34%**
Macro	81.63%	95.24%	99.38%	98.90%	0.709	-0.059	52.23%	53.50%	56.08%

6 Discussion

TACFormer demonstrates that protein-sequence conditioning can steer full-molecule PROTAC generation toward chemistry associated with a requested POI when that target is represented during

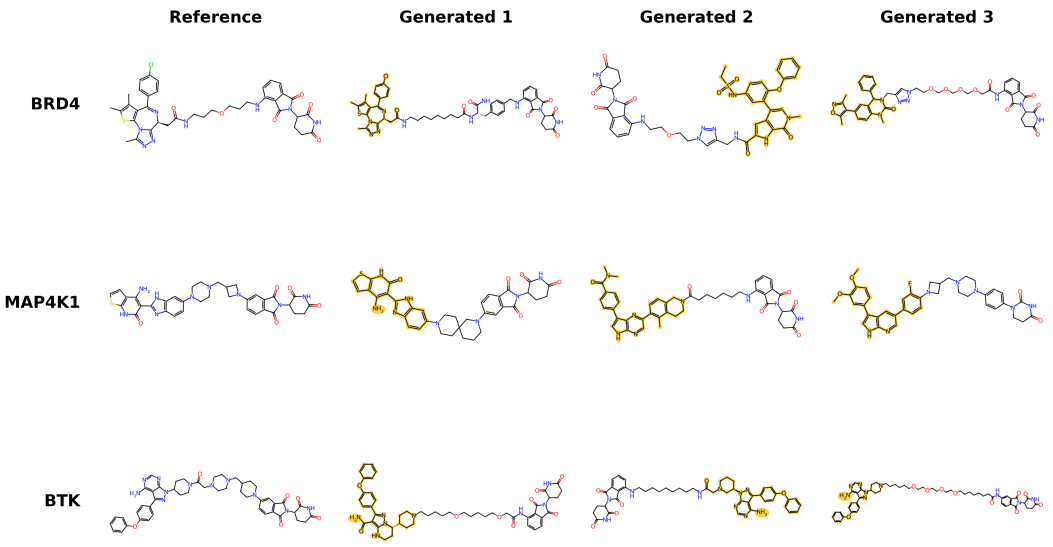

Figure 1: **Generated PROTACs span distinct predicted warhead scaffolds under the same POI condition.** Rows correspond to BRD4, MAP4K1, and BTK, the three evaluated POIs with the highest generated-warhead internal diversity (0.871, 0.841, and 0.826, respectively). The first column shows a reference PROTAC for the corresponding POI, selected as the medoid of its reference set under Morgan–Tanimoto similarity. The remaining columns show novel generated PROTACs carrying distinct Bemis–Murcko scaffolds of the predicted POI-binding warhead. Gold highlighting indicates the POI-binding warhead assigned by PROTAC-Splitter; the remaining structure comprises linker- and E3-ligand-facing chemistry.

training. The consistently positive Δ_{NLL} values across the 20 seen POIs show that the model's likelihood is sensitive to the introduced protein condition, while nearest-neighbor purity shows that this sensitivity is reflected in the generated chemical space. Importantly, however, the strength of this effect varies across targets.

Conditional fidelity depends on the structure of the target-associated chemical space. A single POI may be represented by several distinct medicinal-chemistry series rather than by one compact region of chemical space. TACFormer must represent all of these series using a single fixed protein-derived condition, which makes the conditional learning problem substantially easier for targets associated with chemically concentrated data than for targets spanning multiple chemotypes.

This interpretation is consistent with the observed relationship between conditional fidelity and generated diversity. Δ_{NLL} decreases as the internal diversity of the generated target-specific distribution increases (Appendix D.3), suggesting that the likelihood-level conditioning signal is strongest when the chemistry associated with a POI is comparatively coherent. A similar but weaker trend is observed for nearest-neighbor purity.

The results also highlight an important trade-off between target specificity and generative breadth. METTL14, for example, achieves very high nearest-neighbor purity while producing one of the most concentrated generated distributions. Conversely, more diverse target-conditioned distributions can exhibit weaker local enrichment. Conditional fidelity should therefore not be interpreted independently of uniqueness, novelty, and diversity: high purity can partly reflect concentration around a restricted region of known target-associated chemistry rather than broad exploration of chemically distinct solutions.

The component-level analysis provides complementary evidence that conditional generation is not expressed solely through linker or E3-ligand variation. Generated molecules contain diverse and frequently training-novel predicted POI-binding warheads. Nevertheless, warhead novelty alone does not establish target relevance. A structurally novel warhead may represent useful exploration, but it may also reflect departure from the chemistry associated with the requested target. Warhead-level

novelty is therefore most informative when considered together with whole-molecule conditional-fidelity measures.

Generation becomes less reliable for unseen protein conditions. The two zero-shot targets show a qualitatively different behavior from the seen-target panel. For both WDR5 and DHFR, Δ_{NLL} becomes negative, and nearest-neighbor purity no longer shows the strong local enrichment observed for targets represented during training. These results indicate that the mapping learned during conditional fine-tuning does not directly transfer to these two unseen POIs.

Generation validity also decreases under the unseen conditions, most notably for WDR5. This suggests that protein conditions outside those encountered during training may affect not only where the model places probability mass in chemical space but also the stability of autoregressive generation. However, the present experiments do not establish the mechanism underlying this effect.

Importantly, WDR5 and DHFR were selected as the held-out targets most dissimilar to the training distribution and therefore represent close to the hardest available test of zero-shot transfer in the present dataset. Performance on unseen POIs that are more similar to targets represented during training may therefore lie between the behavior observed for these two targets and that observed for seen POIs. With only two held-out targets, however, the experiment cannot characterize the general limits of zero-shot transfer. What it does establish is that successful conditioning on seen targets does not imply reliable transfer to strongly out-of-distribution protein conditions.

6.1 Limitations

The evaluation measures whether generated molecules occupy chemical regions associated with a requested target; it does not establish that the generated molecules bind to or degrade that target. TACFormer is trained on reported PROTAC structures without explicitly modeling degradation activity such as DC_{50} or $D_{\max}$. Consequently, high conditional fidelity indicates consistency with target-associated PROTAC chemistry rather than predicted biological activity. Prospective application will require additional filtering or modeling of degradation potency, selectivity, synthetic accessibility, and other experimentally relevant properties.

The zero-shot evaluation is limited to WDR5 and DHFR. These targets were intentionally chosen to provide a stringent test by maximizing dissimilarity from the training target distribution. Their failure therefore demonstrates that transfer cannot be assumed under strongly out-of-distribution conditions, but it does not establish how the model behaves for unseen targets at intermediate levels of similarity. A broader zero-shot study spanning a range of distances from the training distribution would be required to characterize this transition systematically.

Finally, each POI is represented by a single sequence-derived embedding. This representation captures protein identity but does not explicitly encode the three-dimensional geometry of the POI surface, the bound warhead pose, the recruited E3 ligase, or the geometry of the ternary complex. These factors are central to productive PROTAC formation and may help distinguish among the multiple chemical series associated with the same target. Incorporating structural and ternary-complex information therefore represents an important direction for extending target-conditioned PROTAC generation.

7 Conclusion

We introduced TACFormer, a transformer-based model for target-conditioned *de novo* generation of complete PROTAC molecules. A molecular prior pretrained on PROTAC-like compounds from ZINC and ChEMBL is subsequently fine-tuned on PROTAC data, with protein-sequence representations providing an explicit POI condition during generation.

Across 20 POIs represented during training, TACFormer generates predominantly valid and novel molecules while exhibiting clear POI-specific chemical steering. Held-out molecules receive higher likelihood under their associated POI condition than under competing conditions, and generated molecules preferentially occupy chemical neighborhoods enriched for the requested target. Component-level analysis further shows that the model can vary the predicted POI-binding chemistry rather than relying exclusively on linker or E3-ligand variation.

The strength of conditioning varies substantially across targets and is related to the heterogeneity of their associated chemical space. Moreover, the same behavior does not transfer reliably to the two strongly out-of-distribution POIs examined here. These findings position protein conditioning as a promising mechanism for target-aware full-molecule PROTAC generation while also identifying generalization beyond observed targets as a central open problem. Future work should combine protein conditioning with structural information and degradation-activity modeling to move from target-associated chemical steering toward the prospective design of active degraders.

AI USE STATEMENT

Generative AI tools were used to edit manuscript prose, propose section organization and figure designs, identify potentially relevant literature. The authors reviewed all AI-assisted material, verified every citation against the primary source, and take responsibility for the final content of this work.

ETHICS STATEMENT

This work raises no concerns involving human subjects, personal data, or animal experimentation. Our PROTAC dataset aggregates three public resources, one of which is community-curated, so entry quality and completeness vary; we describe our normalization and deduplication procedure thoroughly and release the assembled dataset under terms consistent with those sources. Because our metrics assess chemical-space proximity rather than degradation activity, generated molecules should be treated as design hypotheses requiring synthetic and experimental validation, not as predicted degraders. We report the cases where target conditioning fails, including on unseen targets, so that the method's limits are also visible.

REPRODUCIBILITY STATEMENT

Code for TACFormer, downloading and creating the PROTAC and PROTAC-like datasets, and all experiments and analyses presented in this work are available a https://anonymous.4open.science/r/TACFormer-B7AB/.

REFERENCES

Josep Arús-Pous, Simon Viet Johansson, Oleksii Prykhodko, Esben Jannik Bjerrum, Christian Tyrchan, Jean-Louis Reymond, Hongming Chen, and Ola Engkvist. Randomized SMILES strings improve the quality of molecular generative models. *Journal of Cheminformatics*, 11(1):71, 2019.

Miklós Békés, David R Langley, and Craig M Crews. PROTAC targeted protein degraders: the past is prologue. *Nature Reviews Drug discovery*, 21(3):181–200, 2022.

Hengwei Chen and Jürgen Bajorath. Generative design of compounds with desired potency from target protein sequences using a multimodal biochemical language model. *Journal of Cheminformatics*, 16(1):55, 2024.

Yuhao Dai and Fei Zhu. TRL-PROTAC: A pre-trained generator of PROTACs targeting specific proteins optimized by reinforcement learning. *Computer Science and Information Systems*, 21 (4):1293–1320, 2024. doi: 10.2298/CSIS240327039D.

Jingxuan Ge, Shimeng Li, Gaoqi Weng, Huating Wang, Meijing Fang, Huiyong Sun, Yafeng Deng, Chang-Yu Hsieh, Dan Li, and Tingjun Hou. PROTAC-DB 3.0: an updated database of PROTACs with extended pharmacokinetic parameters. *Nucleic Acids Research*, 53(D1):D1510–D1515, 2025.

Yossra Gharbi and Rocío Mercado. A comprehensive review of emerging approaches in machine learning for de novo PROTAC design. *Digital Discovery*, 3(11):2158–2176, 2024.

Jeff Guo, Franziska Knuth, Christian Margreitter, Jon Paul Janet, Kostas Papadopoulos, Ola Engkvist, and Atanas Patronov. Link-INVENT: generative linker design with reinforcement learning. *Digital Discovery*, 2(2):392–408, 2023.

Ilia Igashov, Hannes Stärk, Clément Vignac, Arne Schneuing, Victor Garcia Satorras, Pascal Frossard, Max Welling, Michael Bronstein, and Bruno Correia. Equivariant 3D-conditional diffusion model for molecular linker design. *Nature Machine Intelligence*, 6(4):417–427, 2024.

Fergus Imrie, Anthony R Bradley, Mihaela van der Schaar, and Charlotte M Deane. Deep generative models for 3D linker design. *Journal of Chemical Information and Modeling*, 60(4):1983, 2020.

John J Irwin, Khanh G Tang, Jennifer Young, Chinzorig Dandarchuluun, Benjamin R Wong, Munkhzul Khurelbaatar, Yurii S Moroz, John Mayfield, and Roger A Sayle. Zinc20–a free ultra large-scale chemical database for ligand discovery. *Journal of chemical information and modeling*, 60(12):6065, 2020.

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.

Baiqing Li, Ting Ran, and Hongming Chen. 3D based generative PROTAC linker design with reinforcement learning. *Briefings in Bioinformatics*, 24(5):bbad323, 2023.

Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin, Robert Verkuil, Ori Kabeli, Yaniv Shmueli, et al. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637):1123–1130, 2023.

Nir London and Jaime Prilusky. PROTAC-Pedia: A crowd-sourced database of PROTACs. `https://protacpedia.weizmann.ac.il/`, 2020. Version 1.62. Accessed: 2026-09-24.

Michael Moret, Lukas Friedrich, Francesca Grisoni, Daniel Merk, and Gisbert Schneider. Generative molecular design in low data regimes. *Nature Machine Intelligence*, 2(3):171–180, 2020.

Rebecca M Neeser, Mehmet Akdel, Daniel Kovtun, and Luca Naef. Reinforcement learning-driven linker design via fast attention-based point cloud alignment. *arXiv preprint arXiv:2306.08166*, 2023.

Divya Nori, Connor W. Coley, and Rocío Mercado. De novo PROTAC design using graph-based deep generative models. *arXiv preprint arXiv:2211.02660*, 2022.

Matthew N O'Brien Laramy, Suman Luthra, Matthew F Brown, and Derek W Bartlett. Delivering on the promise of protein degraders. *Nature Reviews Drug Discovery*, 22(5):410–427, 2023.

Andy Pike, Esther CY Lee, Iacovos N Michaelides, Markus Schade, Ankit Sharma, James S Scott, and Abhishek Srivastava. Lessons learned in linking PROTACs from discovery to the clinic. *Nature Reviews Chemistry*, 10(2):117–132, 2026.

Xinran Qin, Yinpeng Zhang, Yajunzi Wang, Yintao Zhang, Jiachen Jing, Yuyuan Zhang, Gaoxiang Xu, Haoping Teng, Tianjun Wang, Lei Fu, et al. TPDdb: the comprehensive database of targeted protein degrader. *Nucleic Acids Research*, 54(D1):D1683–D1691, 2026.

Stefano Ribes, Ranxuan Zhang, Télio Cropsal, Anders Källberg, Christian Tyrchan, Eva Nittinger, and Rocío Mercado. PROTAC-Splitter: a machine learning framework for automated identification of PROTAC substructures. *Journal of Cheminformatics*, 18(1):30. ISSN 1758-2946. doi: 10.1186/s13321-025-01135-9. URL `https://doi.org/10.1186/s13321-025-01135-9`.

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

Jinsong Shao, Qineng Gong, Zeyu Yin, Yu Chen, Yajie Hao, Lei Zhang, Linlin Jiang, Min Yao, Jinlong Li, Fubo Wang, and Li Wang. Language model driven: A PROTAC generation pipeline with dual constraints of structure and property. *arXiv preprint arXiv:2412.09661*, 2024.

David Weininger. SMILES, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of Chemical Information and Computer Sciences*, 28(1):31–36, 1988.

Mingyuan Xu, Chaoming Huang, Li Pang, Qirui Deng, Hao Zhang, Anjie Qiao, Zhiwen Luo, Zhen Wang, Chang-Yu Hsieh, Zhang Zhang, Tie-Gen Chen, Hongming Chen, and Jinping Lei. Design of synthesizable PROTACs through synthesis constrained generative model and reinforcement learning. *JACS Au*, 6(7):4239–4252, 2026. doi: 10.1021/jacsau.6c00732.

Yuyao Yang, Shuangjia Zheng, Shimin Su, Chao Zhao, Jun Xu, and Hongming Chen. SyntaLinker: automatic fragment linking with deep conditional transformer neural networks. *Chemical Science*, 11(31):8312–8322, 2020.

Barbara Zdrazil, Eloy Felix, Fiona Hunter, Emma J Manners, James Blackshaw, Sybilla Corbett, Marleen De Veij, Harris Ioannidis, David Mendez Lopez, Juan F Mosquera, et al. The chembl database in 2023: a drug discovery platform spanning multiple bioactivity data types and time periods. *Nucleic acids research*, 52(D1):D1180–D1192, 2024.

A MODEL AND IMPLEMENTATION DETAILS

This appendix provides implementation and training details for the molecular representation, protein encoding, decoder architecture, conditioning mechanism, and autoregressive training procedure summarized in Section 3.

A.1 SMILES TOKENIZATION AND AUGMENTATION

PROTACs are represented as SMILES and tokenized using the same molecular tokenizer across pretraining, unconditional fine-tuning, and conditional fine-tuning. Beginning-of-sequence (BOS) and end-of-sequence (EOS) tokens are added to each molecule. For a tokenized sequence

$$[\mathrm{BOS}, s_1, \ldots, s_T, \mathrm{EOS}], \tag{9}$$

causal language modeling uses the shifted input–target pair

$$\mathbf{x} = [\mathrm{BOS}, s_1, \ldots, s_T], \qquad \mathbf{y} = [s_1, \ldots, s_T, \mathrm{EOS}]. \tag{10}$$

Sequences are restricted to a maximum length of 320 tokens. Within a minibatch, input sequences are right-padded with the padding token, while the corresponding target positions are assigned the ignore index -100 and therefore do not contribute to the cross-entropy loss.

During training, SMILES randomization is performed on the fly. A randomized SMILES corresponds to an alternative traversal of the same molecular graph and therefore preserves molecular identity while exposing the decoder to different valid serializations of the same structure. Randomization is applied only to the training data; validation examples retain their fixed input representation. In the conditional setting, randomization modifies only the SMILES sequence and never the associated POI condition. Thus, every randomized representation of a molecule remains paired with the same protein embedding.

A.2 PROTEIN SEQUENCE REPRESENTATION

For each POI, the corresponding amino-acid sequence is retrieved from UniProtKB using its normalized UniProt accession. Protein representations are computed using the pretrained `esm2_t33_650M_UR50D` model (Lin et al., 2023), which consists of 33 transformer layers, 650M parameters, and a hidden dimension of 1280.

For a protein sequence

$$p = (a_1, \ldots, a_L), \tag{11}$$

ESM-2 produces contextual sequence representations. The initial ESM special token and terminal token are excluded, leaving one representation for each amino-acid position,

$$\mathbf{H}_p = [\mathbf{h}_1, \ldots, \mathbf{h}_L]^\top \in \mathbb{R}^{L \times 1280}. \tag{12}$$

A single protein-level POI representation is then obtained by mean pooling along the sequence dimension,

$$\mathbf{z}_p = \frac{1}{L} \sum_{i=1}^{L} \mathbf{h}_i \in \mathbb{R}^{1280}. \tag{13}$$

Thus, the hidden feature dimension remains 1280; pooling removes only the variable sequence-length dimension. The resulting vector $\mathbf{z}_p$ is the representation supplied to the conditional molecular model.

ESM-2 is used as a frozen feature extractor. Protein embeddings are precomputed in FP32, stored by UniProt accession, and reused during molecular-model training and generation. ESM-2 parameters are not updated during any stage of conditional fine-tuning.

A.3 LONG-PROTEIN ENCODING

The selected ESM-2 model is applied to windows containing at most 1022 amino-acid residues. Proteins satisfying

$$L \leq 1022 \tag{14}$$

are encoded in a single pass. Longer sequences are divided into overlapping windows of maximum length

$$W = 1022 \tag{15}$$

with an overlap of

$$O = 256 \tag{16}$$

residues. The resulting stride is

$$S = W - O = 766. \tag{17}$$

Each window is processed independently by ESM-2. After excluding the special tokens from each window output, the residue-level representations are reconstructed into a single full-length representation before pooling.

For two adjacent windows, their common overlap is divided at its midpoint. The earlier window supplies representations from the first part of the overlap and the later window supplies representations from the second part. Consequently, although an overlapping residue is processed in both ESM-2 windows and can therefore benefit from the local context of either window, exactly one of its contextual representations is retained in the reconstructed tensor.

The reconstruction is explicitly required to satisfy

$$\mathrm{rows}(\mathbf{H}_p) = L, \tag{18}$$

ensuring that the stitched representation contains exactly one row for every residue in the original sequence.

Mean pooling is applied only after this reconstruction:

$$\mathbf{z}_p = \frac{1}{L} \sum_{i=1}^{L} \mathbf{H}_{p,i}. \tag{19}$$

We therefore do not independently average the representations of individual windows and then average the window-level means. Instead, the complete residue-level representation is first reconstructed and a single mean is subsequently computed over the original protein length.

A.4 DECODER ARCHITECTURE

The molecular model is a decoder-only transformer with eight pre-normalization blocks, hidden dimension $d_{\mathrm{model}} = 512$, eight attention heads, and a feed-forward expansion ratio of four.

For token s_t at sequence position t, the initial representation is obtained by summing learned token and positional embeddings,

$$\mathbf{x}_t^{(0)} = \mathbf{E}_{\mathrm{tok}}(s_t) + \mathbf{E}_{\mathrm{pos}}(t). \tag{20}$$

The resulting sequence representation is propagated through causal transformer blocks. For one attention head, queries, keys, and values are

$$\mathbf{Q} = \mathbf{X}\mathbf{W}_Q, \qquad \mathbf{K} = \mathbf{X}\mathbf{W}_K, \qquad \mathbf{V} = \mathbf{X}\mathbf{W}_V. \tag{21}$$

Causal scaled dot-product attention is

$$\mathrm{Attn}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \mathrm{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_h}} + \mathbf{M} \right) \mathbf{V}, \tag{22}$$

where $d_h = 64$ is the head dimension and $\mathbf{M}$ denotes the causal mask preventing information flow from future tokens.

The attention residual of block ℓ is

$$\mathbf{A}^{(\ell)} = \mathbf{X}^{(\ell)} + \mathrm{MHA}\left(\mathrm{LN}_1\left(\mathbf{X}^{(\ell)}\right)\right). \tag{23}$$

The position-wise feed-forward network expands the hidden representation by a factor of four and uses a GELU activation,

$$\mathrm{MLP}(\mathbf{x}) = \mathbf{W}_2\,\mathrm{GELU}\left(\mathbf{W}_1\mathbf{x} + \mathbf{b}_1\right) + \mathbf{b}_2. \tag{24}$$

In the unconditional model, the block output is therefore

$$\mathbf{X}^{(\ell+1)} = \mathbf{A}^{(\ell)} + \mathrm{MLP}\left(\mathrm{LN}_2\left(\mathbf{A}^{(\ell)}\right)\right). \tag{25}$$

A final layer-normalization operation is applied after the last transformer block, followed by a linear language-model head producing logits over the SMILES vocabulary. The input token embeddings and output projection weights are tied.

A.5 Protein-to-molecule conditioning

Conditional generation retains the same molecular backbone and introduces an adapter connecting the fixed protein representation to the decoder hidden space. The pooled ESM-2 representation

$$\mathbf{z}_p \in \mathbb{R}^{1280} \tag{26}$$

is first normalized and linearly projected:

$$\mathbf{c}_p = \mathbf{W}_c\,\mathrm{LN}(\mathbf{z}_p) + \mathbf{b}_c, \qquad \mathbf{W}_c \in \mathbb{R}^{512\times 1280}, \tag{27}$$

yielding

$$\mathbf{c}_p \in \mathbb{R}^{512}. \tag{28}$$

This projection is computed once for each molecule–POI pair and shared across the decoder blocks.

Each transformer block ℓ contains an independent learnable scalar gate γ_ℓ. Following the self-attention residual, the protein condition is injected according to

$$\widetilde{\mathbf{A}}^{(\ell)} = \mathbf{A}^{(\ell)} + \tanh(\gamma_\ell)\mathbf{c}_p. \tag{29}$$

Because $\mathbf{c}_p$ contains no sequence-position dimension, it is broadcast across all SMILES positions. The subsequent block update is

$$\mathbf{X}^{(\ell+1)} = \widetilde{\mathbf{A}}^{(\ell)} + \mathrm{MLP}\left(\mathrm{LN}_2\left(\widetilde{\mathbf{A}}^{(\ell)}\right)\right). \tag{30}$$

The same projected POI representation is therefore available throughout the decoder, while the independent gates allow different layers to learn different conditioning strengths. The $\tanh$ transformation bounds each scalar modulation while allowing both positive and negative contributions.

A.6 Initialization of the conditional model

The conditional model is initialized from the pretrained unconditional molecular prior rather than from random weights. The token embeddings, positional embeddings, attention layers, feed-forward networks, normalization layers, and language-model head are transferred from the prior checkpoint.

The conditional extension introduces the protein LayerNorm, the $1280 \to 512$ linear projection, and one scalar gate for each of the eight transformer blocks. The protein projection is newly initialized, while every conditioning gate is initialized exactly as

$$\gamma_\ell^{(0)} = 0. \tag{31}$$

Consequently,

$$\tanh(\gamma_\ell^{(0)})\mathbf{c}_p = \mathbf{0}. \tag{32}$$

Immediately after initialization, the added conditioning pathway therefore contributes zero to every block, and the conditional architecture initially reproduces the hidden-state pathway of the pretrained unconditional decoder. During conditional fine-tuning, gradients update the gates, projection, and molecular decoder parameters, allowing the protein signal to influence generation progressively without introducing an abrupt randomly initialized perturbation at initialization.

Conceptually, conditional fine-tuning transforms the pretrained molecular distribution

$$P_{\theta_0}(\mathbf{s}) \tag{33}$$

into

$$P_\theta(\mathbf{s} \mid \mathbf{z}_p), \tag{34}$$

using the pretrained molecular parameters θ_0 as the starting point.

A.7 Causal language-model objective

For conditional training, each molecular example is paired with the precomputed representation of its associated POI. For minibatch size B, the protein vectors are stacked as

$$\mathbf{Z} = [\mathbf{z}_{p_1}, \ldots, \mathbf{z}_{p_B}]^\top \in \mathbb{R}^{B \times 1280}. \tag{35}$$

The nth protein representation therefore remains aligned with the nth SMILES sequence throughout the batch.

Training uses teacher-forced causal language modeling. Given a training pair $(\mathbf{s}^{(n)}, p_n)$, the objective is the token-level negative log-likelihood

$$\mathcal{L}_{\mathrm{LM}} = -\frac{1}{N} \sum_{n=1}^{N} \sum_{t=1}^{T_n} \log P_\theta \left(s_t^{(n)} \mid s_{<t}^{(n)}, \mathbf{z}_{p_n} \right). \tag{36}$$

At position t, the decoder therefore receives the preceding ground-truth SMILES tokens together with the fixed protein representation and predicts the next ground-truth token. Cross-entropy is computed over the decoder logits, with padded target positions ignored.

The ESM-2 parameters are not part of this optimization. Conditional learning occurs in the protein adapter, the layer-specific gates, and the molecular decoder itself.

A.8 Optimization details

Conditional fine-tuning uses a per-device training batch size of 32 and validation batch size of 32 with no additional gradient accumulation. The learning rate is 1×10^{-4}, with weight decay 0.01, gradient clipping at a maximum norm of 1.0, and a cosine learning-rate schedule with a warm-up ratio of 0.05. Training is performed using bfloat16 mixed precision.

A.9 Autoregressive generation

At inference time, generation starts from a single BOS token. In the conditional setting, the selected POI embedding is retrieved from the precomputed lookup and supplied to the model for the complete decoding trajectory. The same fixed $\mathbf{z}_p$ is therefore used at every autoregressive step.

For decoder logits $\mathbf{o}_t$ at step t, sampling probabilities are obtained as

$$P_\theta \left(s_t \mid s_{<t}, \mathbf{z}_p \right) = \mathrm{softmax} \left(\frac{\mathbf{o}_t}{\tau} \right), \tag{37}$$

where τ denotes the temperature.

Generation supports top-k and nucleus (top-p) sampling. For nucleus sampling, tokens are ranked by probability and sampling is restricted to the smallest set whose cumulative probability reaches the specified threshold. PAD, BOS, and unknown tokens are prohibited during generation, whereas EOS remains available so that decoding can terminate naturally.

When no explicit generation-length limit is supplied, the maximum number of newly generated tokens is set to one less than the model context length, corresponding to 319 tokens when generation begins from a single BOS token. Generation terminates earlier when EOS is sampled.

The complete conditional generation pathway is therefore

$$\text{protein sequence} \rightarrow \text{frozen ESM-2} \rightarrow \mathbf{H}_p \rightarrow \mathbf{z}_p \rightarrow \mathbf{c}_p \rightarrow \text{gated decoder} \rightarrow \text{SMILES}. \tag{38}$$

B Additional Experimental Details

B.1 PROTAC-like pretraining corpus

The PROTAC-like pretraining corpus described in Section 4 was constructed using physicochemical and structural filters intended to enrich for large bifunctional molecules. Compounds were required to have molecular weight between 650 and 1200 Da, TPSA between 120 and 320 Å^2, 12–45 rotatable bonds, 8–24 hydrogen-bond acceptors, and 0–8 hydrogen-bond donors.

To further enrich for bifunctional architectures, molecules were required to contain at least two distinct ring systems. The two largest ring systems were identified and the minimum topological distance between them was restricted to 8–60 bonds. Molecules satisfying all criteria were canonicalized before merging and deduplication across ZINC20 and ChEMBL36.

B.2 Dataset composition and component-level chemistry

Figure 2 provides additional characterization of the conditional PROTAC dataset. The 20 most represented POIs account for 70.7% of all molecules (Figure 2a,b), illustrating the strong concentration of available PROTAC data around a relatively small number of targets.

E3-ligase usage is similarly concentrated. CRBN and VHL dominate the dataset globally (Figure 2c), while their relative contributions vary considerably across the individual top-20 POIs (Figure 2d).

To examine the chemistry of the two binding ends separately, PROTAC-Splitter was applied to the 21,172 unique structures associated with the top-20 POIs. Of these, 20,084 (94.9%) passed component-level quality control. Bemis–Murcko scaffolds were extracted from the predicted POI-binding warhead and E3-ligand components and deduplicated within each POI. Pairwise similarity between POIs was then quantified using nearest-scaffold Tanimoto similarity.

As shown in Figure 2e, E3-ligand scaffold repertoires are substantially more similar across POIs than POI-warhead scaffold repertoires. This motivates controlling for E3 identity when assessing whether whole-PROTAC chemical organization reflects POI identity, as done in the E3-stratified analyses of Figure 3.

B.3 Conditional data splitting

For the seen-target conditional experiment, the 20 most represented POIs formed the evaluation panel. Molecule–POI associations for these targets were split approximately 80/10/10 into training, validation, and test sets. Associations for the remaining POIs were retained only in the training set, exposing the model to a broader range of protein conditions while restricting held-out evaluation to targets with sufficient reference chemistry.

WDR5 and DHFR were reserved exclusively for zero-shot evaluation and removed before constructing the seen-target partitions. Their associated PROTAC chemistry was therefore excluded from conditional training, validation, and seen-target testing. At zero-shot inference, the same trained conditional model was conditioned on the pretrained protein representations of WDR5 or DHFR without additional fine-tuning.

B.4 POI-associated chemical-space characterization

Chemical-space analyses use 2048-bit Morgan fingerprints with radius 2 and Tanimoto similarity. For global comparisons, same-POI and different-POI pairs were sampled by first selecting POIs and then selecting molecules, reducing domination by the most highly represented targets.

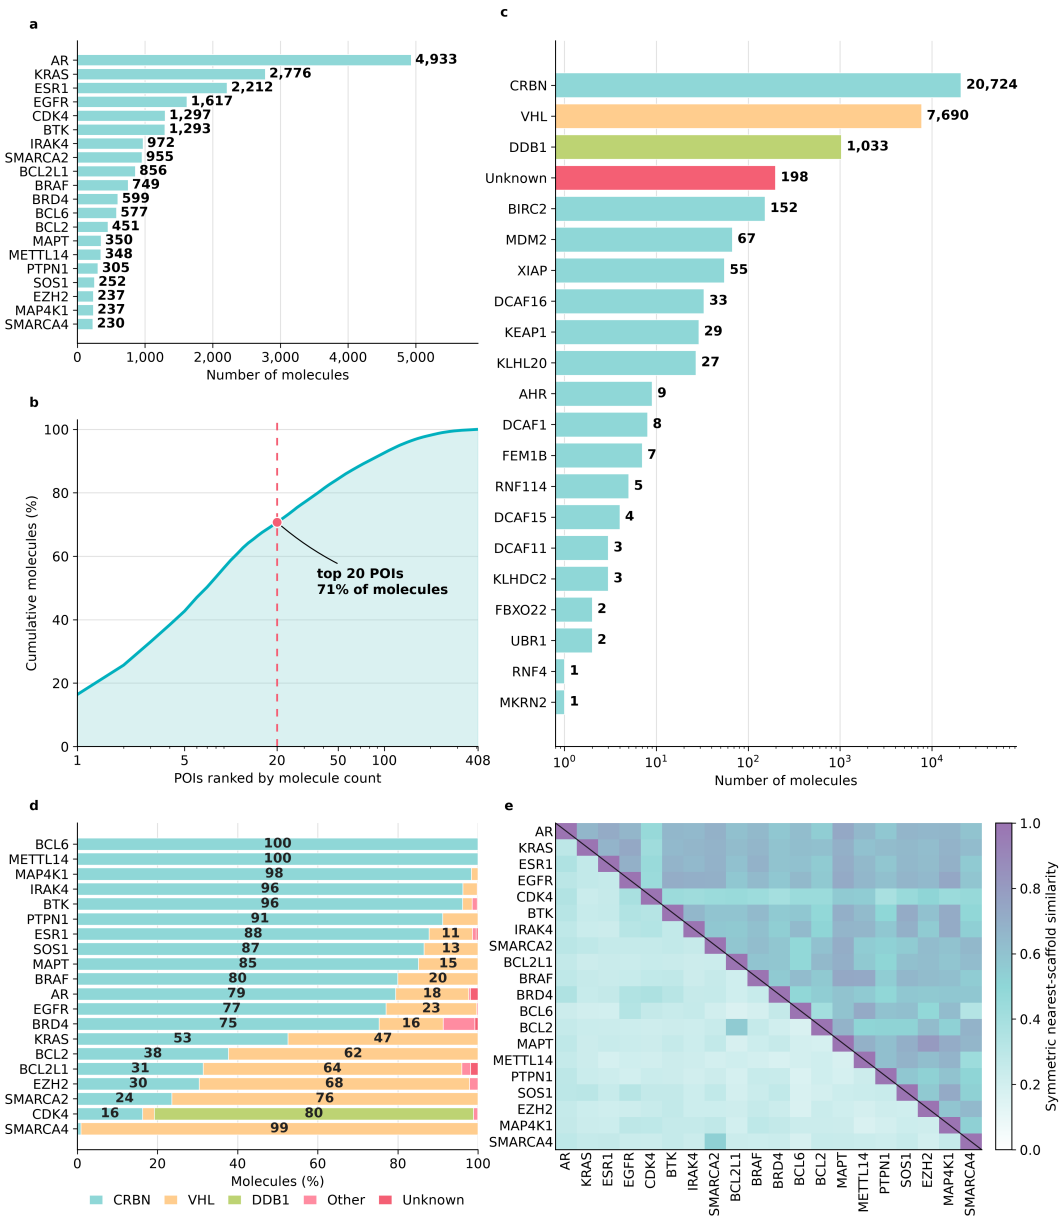

Figure 2: **Dataset composition and component-level scaffold similarity.** **(a)** Number of PROTAC associations for the 20 most represented POIs. **(b)** Cumulative fraction of the conditional dataset covered by POIs ranked by number of associations. **(c)** Distribution of E3 ligases across the complete conditional dataset. **(d)** E3-ligase composition of the top 20 POIs. **(e)** Pairwise similarity of predicted POI-warhead (lower triangle) and E3-ligand (upper triangle) scaffold repertoires across the top 20 POIs, measured by nearest-scaffold Tanimoto similarity.

Because shared E3 chemistry can create whole-molecule similarity independently of the POI-binding end, pairwise comparisons were additionally performed separately within CRBN and VHL backgrounds.

Local target organization was assessed using nearest-neighbor POI purity for neighborhood sizes $k = 1, \ldots, 30$. The observed purity curves were compared with a null distribution obtained from 300 permutations of POI labels within each E3 group while leaving molecular fingerprints and neighbor relationships unchanged.

For visualization, 300 molecules were sampled from each of the 10 most represented POIs. Morgan fingerprints were first reduced to 50 principal components and subsequently embedded using t-SNE. The same embedding coordinates were used when coloring molecules by POI and E3 identity.

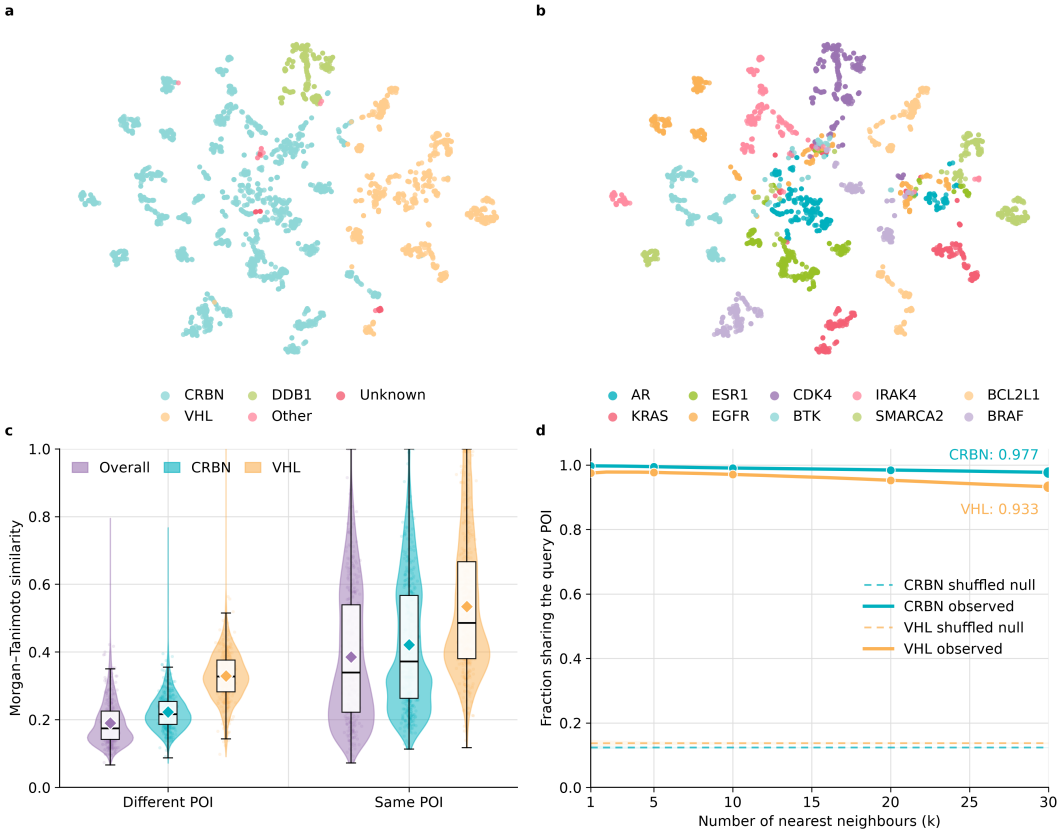

Figure 3: **POI-associated PROTAC chemistry is globally heterogeneous but locally organized.** All panels use 2048-bit Morgan fingerprints with radius 2, matching the settings used for the generation metrics. **(a,b)** The same two-dimensional embedding of PROTAC chemical space, colored by **(a)** E3-ligase and **(b)** POI identity; both panels share coordinates and differ only in coloring. The embedding contains 3,000 molecules, 300 sampled from each of the 10 most represented POIs, reduced to 50 principal components and embedded by t-SNE. E3 identity organizes the space broadly **(a)**, whereas POIs form enriched local regions and individual POIs can occupy several separated regions **(b)**. **(c)** Morgan–Tanimoto similarity for same-POI and different-POI PROTAC pairs over the top-20 evaluation panel. Same-POI pairs are more similar overall (0.385 versus 0.190) and within each major E3 background separately (CRBN: 0.421 versus 0.223; VHL: 0.535 versus 0.329), so the signal is not an artifact of E3 chemistry. **(d)** POI purity within fixed E3 groups: for each CRBN- or VHL-associated query molecule, the fraction of its $k = 1$–30 nearest chemical neighbors sharing its POI. Solid lines show observed purity, dashed lines the null obtained by shuffling POI labels within each E3 group over 300 permutations. At $k = 30$, observed purity is 0.977 (CRBN) and 0.933 (VHL). Thus, although molecules associated with the same POI can show moderate global similarity because they may belong to different chemotypes, molecules that are locally close in chemical space are strongly enriched for the same POI.

B.5 Generation protocol

The same stochastic decoding settings were used for the final unconditional and conditional generation experiments. Generation begins from a BOS token and uses temperature $\tau = 1.0$ with nucleus sampling at $p = 0.9$ and no top-k restriction. Generation uses random seed 42 and batch size 32.

PAD, BOS, and UNK tokens are excluded from the sampling distribution, whereas EOS remains available to terminate a sequence. When `max_new_tokens` is unspecified, the implementation sets the generation limit to one token fewer than the model context length. With a context length of 320, this corresponds to at most 319 newly generated tokens following BOS.

The generation procedure outputs raw decoded SMILES only. Chemical parsing, canonicalization, deduplication, and metric calculation are performed subsequently during evaluation; no chemistry-based post-hoc filtering is applied to the generated strings before evaluation.

For conditional generation, the selected POI representation remains fixed throughout the complete autoregressive trajectory of each molecule. The same decoding procedure is used for seen and unseen POIs.

B.6 Generation-quality metric implementation

Generated SMILES are parsed using RDKit `Chem.MolFromSmiles`. A molecule is counted as valid when parsing and standard RDKit sanitization succeed; no additional connected-component requirement is imposed. Valid structures are converted to canonical isomeric SMILES, retaining encoded stereochemical information.

For conditional evaluation, occurrence counts and deduplication are recomputed independently within each requested POI. Repeated generation of the same canonical structure under one condition therefore reduces that POI's uniqueness, whereas the same structure generated under two different POI conditions can contribute once to each condition.

Whole-molecule novelty is evaluated against the complete corresponding training chemistry rather than against a POI-specific subset. Thus, for the conditional model, a generated structure is considered non-novel if the same canonical molecule appeared anywhere in the conditional training split, even if it was associated with a different POI.

Bemis–Murcko scaffolds are extracted from valid molecules and converted to canonical isomeric scaffold SMILES. Scaffold novelty is evaluated against the complete scaffold repertoire observed in the corresponding training data.

Internal diversity is calculated on unique valid molecules using 2048-bit Morgan fingerprints with radius 2 and mean pairwise Tanimoto distance. When more than 1,000 unique valid molecules are available, 1,000 molecules are sampled with seed 42 before pairwise similarity calculation.

Physicochemical descriptors are calculated over unique valid molecules.

B.7 Physicochemical comparison of PROTAC-like, PROTAC, and generated molecules

Figure 4 compares the physicochemical distributions of the PROTAC-like pretraining corpus, real PROTACs, and molecules generated by the conditional model. Within each property, all datasets are shown using common histogram bins and normalized densities, and dashed vertical lines indicate the corresponding distribution means.

The PROTAC-like corpus overlaps substantially with the real PROTAC data across the displayed properties, while retaining expected differences arising from its construction by filtering general small-molecule databases. The generated distribution likewise overlaps broadly with real PROTAC chemistry across the examined descriptors, providing a complementary view to the scalar generation-quality metrics reported in the main text.

B.8 Conditional-fidelity reference construction

For each conditional evaluation, the held-out test data are restricted to the same POI panel used in the corresponding generation experiment. Canonical test structures are represented once in the chemical reference pool while retaining all valid POI annotations associated with each molecule.

For the likelihood-based evaluation, each unique test molecule is scored under every POI condition in the corresponding panel. Sequence NLL is normalized by the number of non-padding prediction

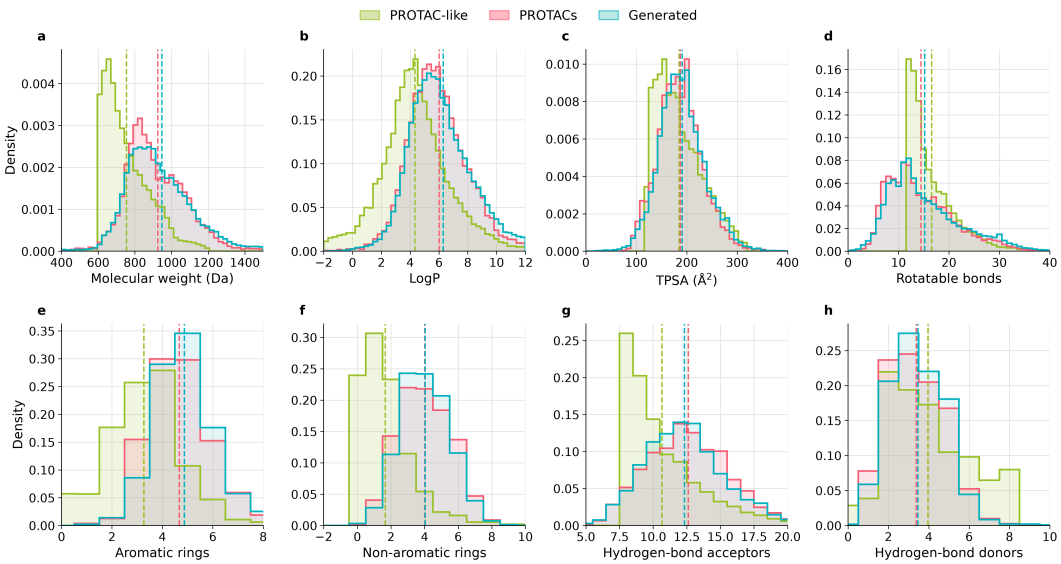

Figure 4: **Physicochemical distributions of PROTAC-like, real PROTAC, and generated molecules.** Normalized density distributions for PROTAC-like structures, real PROTACs, and molecules generated by the conditional model. Dashed vertical lines indicate distribution means. **(a)** Molecular weight (MW). **(b)** LogP. **(c)** Topological polar surface area (TPSA). **(d)** Number of rotatable bonds. **(e)** Number of aromatic rings. **(f)** Number of non-aromatic rings. **(g)** Number of hydrogen-bond acceptors (HBA). **(h)** Number of hydrogen-bond donors (HBD). For visualization, the x-axis ranges are truncated to emphasize the regions containing the bulk of each distribution; molecules outside the displayed ranges remain included in the underlying datasets and summary statistics.

tokens. When a molecule has multiple legitimate POI annotations, all such annotations are excluded from the competing-condition set rather than being treated as incorrect targets.

For nearest-neighbor fidelity, generated queries are restricted to molecules that are valid, absent from the complete conditional training chemistry, and unique within the requested POI condition. There is no fixed cap on the number of final queries. Query and reference molecules are represented using 2048-bit Morgan fingerprints with radius 2.

If a training-novel generated molecule exactly matches a molecule in the held-out test partition, the resulting Tanimoto similarity of 1 is retained rather than removed. Such a match represents agreement with unseen test chemistry and is therefore a valid conditional-fidelity signal.

Nearest-neighbor purity is evaluated for $k \in \{5, 10, 20\}$ using the same multi-label POI annotations retained in the held-out reference pool.

B.9 POI-WARHEAD PROCESSING

POI-binding components are obtained from PROTAC-Splitter decompositions. The warhead analysis begins from unique valid generated PROTACs and is carried out independently for each requested POI. Training reference warheads are constructed using the same decomposition procedure on training molecules associated with that POI.

PROTAC-Splitter represents the linker attachment using mapped dummy atoms. For the reported metrics, the POI-side attachment atom `[*:1]` is replaced with hydrogen to recover a parent-warhead representation. The resulting molecule is sanitized, explicit hydrogens are removed, and the structure is converted to canonical SMILES. Identical processing is applied to generated and training warheads before comparison.

Warhead novelty is evaluated against the processed training warheads for the *same POI*, rather than against the complete training chemistry. Warhead internal diversity is computed over all unique

successfully processed generated warheads using 2048-bit Morgan fingerprints with radius 2; unlike the whole-molecule diversity calculation, no 1,000-molecule subsampling cap is applied.

Because component-level evaluation depends on a successful decomposition, splitter-prediction coverage and warhead-processing coverage are additionally recorded for each POI. These quantities indicate how much of the unique valid generated set contributes to the reported warhead-level metrics.

C Generation Dynamics During Training

Figure 5 summarizes the evolution of generation quality during training of the molecular prior, unconditional model, and conditional model. Validity generally increases or remains high as training progresses, indicating that the models learn and retain the syntactic constraints required to generate valid SMILES. In contrast, novelty gradually decreases as training proceeds, reflecting increasing agreement with the chemistry represented in the corresponding training distributions. Internal diversity remains comparatively stable for the molecular prior and unconditional model, whereas the conditional model exhibits a more pronounced decrease during early fine-tuning before approaching a stable regime.

For the conditional model, generation quality was additionally monitored separately for each of the 20 evaluation POIs (Figure **??**). Validity approaches one across training, while uniqueness and novelty decrease gradually but remain substantial at later epochs. Internal diversity within individual POI conditions decreases during the early stages of conditional fine-tuning and subsequently stabilizes. The broad variation across POIs, represented by the shaded ± 1 standard-deviation intervals, indicates that the evolution of generated chemistry differs substantially between targets.

D Additional Results

D.1 POI-warhead generation analysis

Table 3 provides the complete target-wise warhead analysis underlying the component-level result summarized in Section 5.2. Warhead uniqueness, internal diversity, and novelty are computed from the predicted POI-binding components obtained using PROTAC-Splitter, following the processing procedure described in Appendix B.9.

The target-wise results reveal substantial variation in how diversity is distributed at the POI-binding end. BRD4 and MAP4K1 exhibit the highest warhead internal diversity, whereas BCL6 and METTL14 generate comparatively restricted sets of distinct warheads. Whole-molecule and warhead-level uniqueness can also differ markedly. For example, BCL6 has substantially lower warhead uniqueness than whole-molecule uniqueness, indicating that variation elsewhere in the PROTAC can coexist with repeated use of a smaller POI-binding repertoire.

Warhead novelty should be interpreted jointly with conditional-fidelity metrics. Targets can exhibit high warhead novelty even when their generated whole molecules show limited enrichment around held-out chemistry for the requested POI. Thus, novelty at the binding end is evidence of structural variation relative to known target-associated warheads, but not by itself evidence that the new warheads preserve target relevance.

D.2 Chemical-space organization of generated PROTACs

Figure 7 visualizes the organization of the generated molecules across the 20 seen POI conditions. All generated structures are represented using 2048-bit Morgan fingerprints with radius 2 and embedded jointly in a common two-dimensional t-SNE space.

The generated distributions exhibit target-associated regional structure while remaining heterogeneous within individual POI conditions. Several conditions occupy multiple regions rather than forming single compact clusters, consistent with the multimodal nature of the corresponding PROTAC chemical space.

Because this embedding contains only generated molecules, it should be interpreted as a visualization of differences among the conditional output distributions. It does not measure agreement with

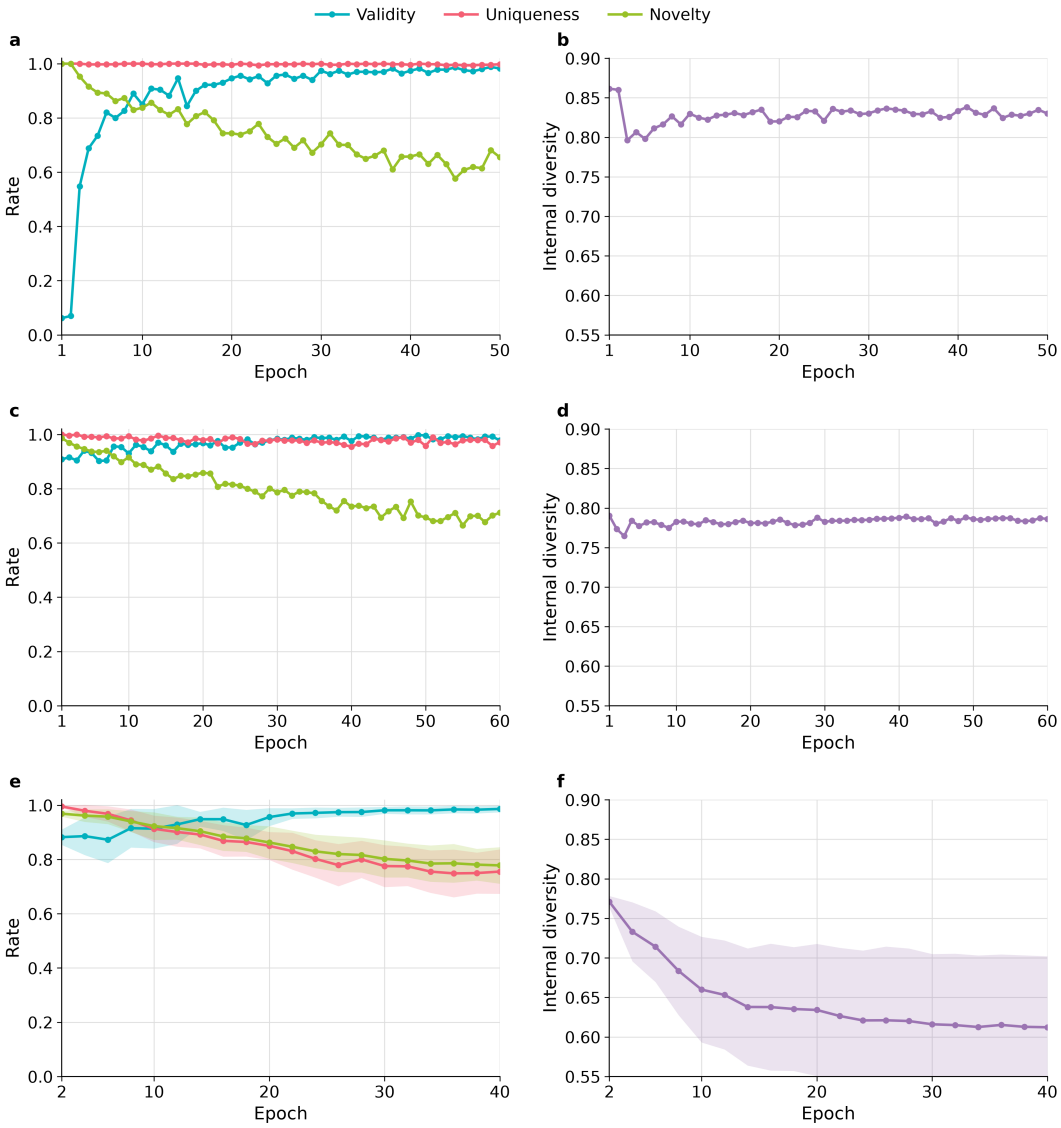

Figure 5: **Generation quality during training of the molecular prior, unconditional model, and conditional model.** Training-time generation metrics are shown for the molecular prior (**a**, **b**), the unconditional model fine-tuned on PROTAC data (**c**, **d**), and the POI-conditioned model (**e**, **f**). **(a,c,e)** Validity, uniqueness among valid molecules, and novelty among unique valid molecules across training epochs. **(b,d,f)** Internal diversity among unique valid molecules. For the prior and unconditional models, curves correspond to generation metrics computed at each monitored epoch. For the conditional model, metrics are computed separately for each monitored POI; solid lines show the mean across POIs and shaded regions indicate ± 1 standard deviation.

real target-associated chemistry; that relationship is quantified separately by the held-out nearest-neighbor analysis reported in Section 5.2.

D.3 Relationship between conditional fidelity and generated diversity

To further characterize the target-to-target heterogeneity observed in Section 5.2, we examined the relationship between internal diversity and the two conditional-fidelity measures across the 20 seen POIs.

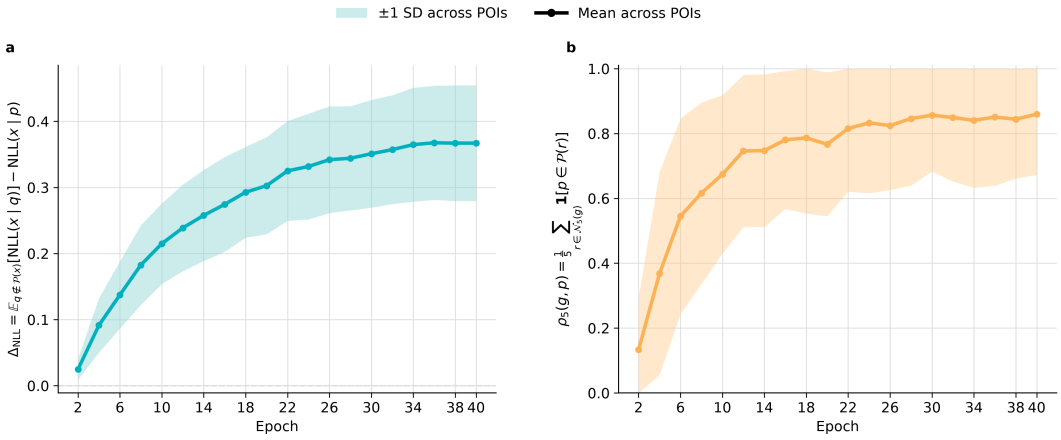

Figure 6: **Use of the POI condition during training. (a)** Δ_{NLL} and **(b)** 5-NN purity ρ_5, evaluated at monitored epochs to track whether the model increasingly uses the supplied POI condition. Solid lines show the mean across the 20 monitored POIs and shaded regions indicate ± 1 standard deviation across POIs. Both metrics use a fixed panel of validation molecules held constant throughout training; at each monitored epoch, 500 molecules are generated per POI, of which up to 25 unique valid molecules per POI serve as queries for ρ_5 against that panel. Mean Δ_{NLL} increases from 0.025 at epoch 2 to 0.367 at epoch 40, and mean ρ_5 from 0.133 to 0.860 over the same interval, indicating that both molecular likelihood **(a)** and the chemical space occupied by generated molecules **(b)** become increasingly dependent on the supplied POI condition.

Table 3: **Warhead-level generation metrics across the 20 seen evaluation POIs.** Warhead uniqueness, internal diversity, and target-specific novelty are calculated on predicted POI-binding warheads obtained using PROTAC-Splitter. Novelty is evaluated against processed training warheads associated with the same POI. The final row reports the unweighted mean across targets. Best per-target values in each column are shown in bold.

POI	Warhead uniqueness	Warhead diversity	Warhead novelty
AR	21.27%	0.815	63.90%
ESR1	20.80%	0.781	70.48%
CDK4	21.06%	0.780	76.67%
KRAS	18.61%	0.742	82.63%
EZH2	29.01%	0.794	97.42%
BTK	33.37%	0.826	87.76%
BRD4	31.54%	**0.871**	95.70%
BRAF	18.77%	0.748	92.13%
SMARCA4	34.50%	0.776	96.97%
MAP4K1	39.00%	0.841	**97.95%**
IRAK4	32.10%	0.821	93.44%
EGFR	28.69%	0.803	86.34%
BCL2L1	22.13%	0.764	87.64%
SMARCA2	33.73%	0.781	89.96%
METTL14	15.13%	0.578	91.30%
PTPN1	21.65%	0.776	94.60%
BCL2	**48.97%**	0.664	95.80%
SOS1	30.68%	0.801	96.83%
BCL6	11.51%	0.809	91.73%
MAPT	21.66%	0.756	94.36%
Macro	26.71%	0.776	89.18%

The fitted relationship between internal diversity and Δ_{NLL} explains 36% of the across-target variation ($R^2 = 0.36$) and is significant at $\alpha = 0.01$. In contrast, the relationship with 5-NN purity is weaker ($R^2 = 0.13$) and is not statistically significant at the present sample size.

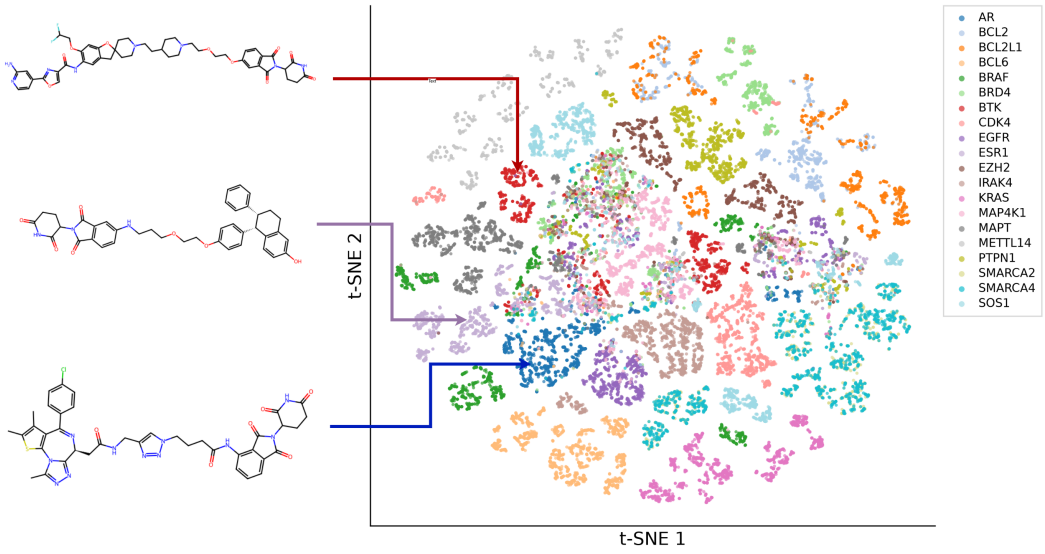

Figure 7: **Chemical-space organization of generated PROTACs across POI conditions.** Two-dimensional t-SNE embedding of molecules generated under the 20 seen POI conditions, colored by conditioning POI. Molecules are represented using 2048-bit Morgan fingerprints with radius 2 and embedded in a single common fit. The example structures shown alongside the embedding are novel generated PROTACs, with arrows indicating the POI condition under which each molecule was generated. The embedding contains only generated molecules and therefore visualizes differences among conditional output distributions rather than proximity to held-out reference chemistry.

These two relationships need not coincide because they quantify different aspects of conditional behavior. Δ_{NLL} measures how strongly the model's likelihood changes with the supplied protein condition on held-out molecules, whereas nearest-neighbor purity measures where generated structures lie relative to held-out target-associated chemistry. The correlations are descriptive and do not establish that diversity itself causes differences in conditional fidelity.

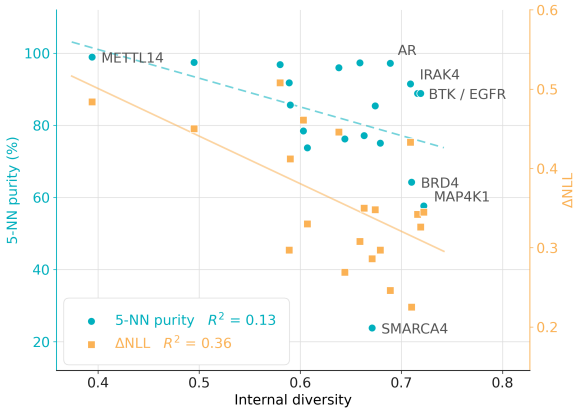

Figure 8: **Relationship between generated diversity and conditional-fidelity measures across seen POIs.** Each point represents one of the 20 seen evaluation POIs. Teal circles, referenced to the left axis, show 5-NN POI purity, while orange squares, referenced to the right axis, show Δ_{NLL}. Both are plotted against internal diversity. Lines correspond to ordinary least-squares fits. The relationship with 5-NN purity is not statistically significant ($R^2 = 0.13$), whereas the relationship with Δ_{NLL} is significant ($R^2 = 0.36$, $\alpha = 0.01$). The two vertical axes use independent scales; only trends within each series are interpretable. Correlations are computed over $n = 20$ POIs without correction for multiple comparisons.