MOLDERIVEX: MULTI-OBJECTIVE OPTIMIZATION VIA ADAPTIVE STRUCTURAL INHERITANCE

Anonymous authors
Paper under double-blind review

ABSTRACT

Multi-objective molecular optimization requires translating competing property objectives into effective structural edits. A key challenge is to identify and reuse substructures that support different properties within the current molecule, while updating this structural guidance as optimization proceeds. We introduce MolDeriveX, a graph optimization framework based on adaptive structural inheritance. It identifies property-specific anchors by combining isolated-fragment scores with changes in parent-level property scores after fragment deletion. These anchors guide Pareto-based graph optimization through soft retention penalties, favoring the retention of property-supporting regions while keeping them editable. Selected high-quality descendants become new structural references, with anchors recomputed to guide subsequent search. Across five multi-objective evaluated tasks, MolDeriveX achieves the highest hypervolume on four and the highest Top-100 joint score on all four target activity tasks. In a challenging, real-world-motivated antibiotic derivatization case study involving three single-parent optimization tasks, MolDeriveX achieves the highest hypervolume for predicted antibacterial activity and parent-molecule similarity among the evaluated methods.

1 INTRODUCTION

Drug discovery requires molecules to combine biological activity, drug-likeness, and synthetic accessibility. These properties can conflict: a structural modification that improves one may weaken another (Zeng et al., 2022; Torren-Peraire et al., 2026). In molecular derivatization, an existing parent also provides structural information worth reusing (Makara et al., 2021). The challenge is to translate competing property objectives into decisions about which regions to preserve, modify, and reconsider during search.

Direct-search methods guide molecular transformations through different mechanisms. MARS learns edit proposals (Xie et al., 2021); GA+DL augments genetic search with a neural discriminator (Nigam et al., 2020); MolSearch searches transformation rules (Sun et al., 2022); and InversionGNN coordinates property gradients during graph inversion (Niu et al., 2025). Pareto evolutionary search also balances population-level objectives (Verhellen, 2022), while rationale-based generation provides property-related fragments (Jin et al., 2020b). However, selecting a promising molecule does not identify which regions support its properties, and isolated fragment scores do not establish their contribution within that parent. Identifying what to preserve from the starting molecules therefore remains a challenge for subsequent multi-objective search (Figure 1A).

Other approaches exploit learned representations, policies, or exemplars. Continuous molecular encodings support latent search (Gómez-Bombarelli et al., 2018); Hgraph provides hierarchical motif-based generation (Jin et al., 2020a), and MLPS combines latent Pareto-set learning with local Bayesian optimization (Liu et al., 2025). HN-GFN learns preference-conditioned sampling (Zhu et al., 2023), while RetMol retrieves exemplars and iteratively updates molecules and its retrieval database (Wang et al., 2023). Such guidance supports generation and objective trade-offs. However, it does not by itself specify which regions should be retained during a particular edit. Fixed structural restrictions can exclude beneficial changes, while unrestricted edits can discard support for existing properties. Balancing retention and revision as the molecule evolves thus remains a challenge beyond objective-level guidance (Figure 1B).

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

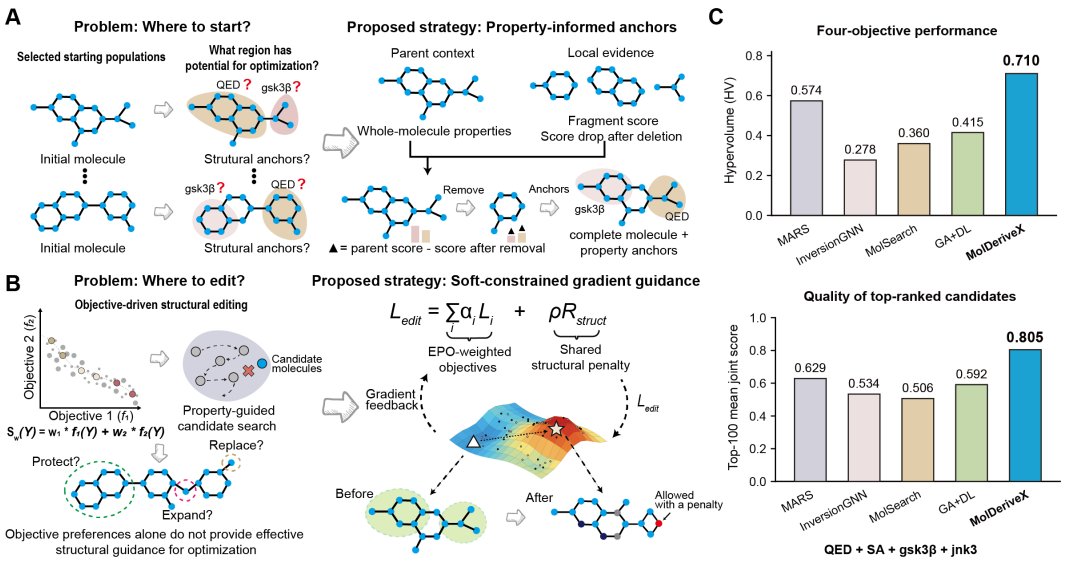


Figure 1: Motivation and strategy of adaptive structural inheritance. (A) Selecting property-specific structural references from the supplied molecules using fragment scores and parent deletion effects. (B) Connecting objective guidance to soft retention during structural editing. (C) Four-objective results against the illustrated direct-search baselines; Tables 1 and 2 give the complete comparisons.

We propose **MolDeriveX**, a multi-objective graph optimization framework based on *adaptive structural inheritance*. It selects property-specific anchors using fragment scores and changes in parent scores after fragment deletion. Matched anchors define soft penalties on graph edits under Pareto guidance. Lineage history guides exploration, and qualifying descendants can become new roots with refreshed anchors. These stages connect structural attribution, editable retention, and reference updates within one search cycle (Figure 2). Our contributions are threefold:

- **Property-specific structural inheritance.** We formulate reusable structural guidance as a revisable state carried by molecular lineages, with anchors selected from both fragment scores and their deletion effects in the parent.
- **Coupled retention and exploration.** We combine soft anchor costs with Pareto-guided graph editing, lineage-aware rescue, and descendant promotion, allowing both molecular structures and their structural references to evolve.
- **Evaluation across tasks and applications.** Comparisons with eight baselines across five multi-objective benchmark tasks show leading performance. Ancestor analyses and ablations assess structural exploration; one real-world antibiotic derivatization case study comprising three single-parent settings evaluates predicted activity–similarity trade-offs.

2 Related Work

Direct search through molecular transformations. Genetic algorithms optimize discrete molecular structures through mutation and recombination (Jensen, 2019; Nigam et al., 2020), including Pareto selection (Verhellen, 2022). MIMOSA uses property- and similarity-constrained sampling (Fu et al., 2021); MARS learns edit proposals (Xie et al., 2021); MolSearch searches over transformation rules (Sun et al., 2022). DST relaxes structural edits for gradient optimization (Fu et al., 2022), and InversionGNN applies Pareto gradients to multi-property graph inversion (Niu et al., 2025). However, objective-level guidance and candidate selection alone do not specify which property-related regions should be retained across successive edits.

Optimization through latent representations. Continuous encodings support property-guided search (Gómez-Bombarelli et al., 2018); JT-VAE and hierarchical graph models decode substructures and motifs (Jin et al., 2018; 2020a). Constrained Bayesian optimization addresses invalid latent

proposals (Griffiths & Hernández-Lobato, 2020). MLPS combines local Bayesian models and adaptive trust regions with a learned latent Pareto set (Liu et al., 2025). However, a property-directed move in latent space does not by itself specify which local molecular regions will be retained after decoding.

Generation through learned policies and exemplars. GCPN and MolDQN learn construction or modification policies (You et al., 2018; Zhou et al., 2019). HN-GFN learns preference-conditioned sampling for multi-objective Bayesian optimization (Zhu et al., 2023). RetMol retrieves exemplars and iteratively updates molecules and its retrieval database (Wang et al., 2023). Rationale-based generation extracts property-related fragments and learns to complete their combinations (Jin et al., 2020b). However, isolated fragment scores and whole-molecule exemplars do not directly establish a selected region's contribution within the current parent. Their relevance can also change as successive edits alter the molecular context and the trade-offs among properties.

3 PRELIMINARIES

Starting from a supplied set of molecules, we seek improved candidates that express different objective trade-offs. We first define objective-space comparison and the molecular edit space; Section 4 connects them through structural anchors and lineage updates.

3.1 PARETO OPTIMIZATION

Let $\mathcal{G}$ denote the space of valid molecules and $\mathbf{f}(G) = [f_1(G), \ldots, f_m(G)]$ the property vector of a molecule $G \in \mathcal{G}$. We use normalized rewards $f_i(G) \in [0, 1]$, all oriented toward maximization. Multi-objective molecular optimization seeks solutions to

$$\max_{G \in \mathcal{G}} \mathbf{f}(G). \tag{1}$$

A molecule G_a dominates G_b, denoted by $G_a \succ G_b$, if

$$\begin{aligned} &f_i(G_a) \geq f_i(G_b), \quad \forall i \in \{1, \ldots, m\}, \\ &f_j(G_a) > f_j(G_b), \quad \text{for some } j. \end{aligned} \tag{2}$$

The Pareto-optimal set contains molecules that are not dominated by any feasible molecule; its image in objective space forms the Pareto front (Censor, 1977). Under a finite evaluation budget, the search instead maintains non-dominated candidates among the evaluated molecules, which approximate rather than certify the global Pareto front.

A preference vector $\mathbf{r} \in \mathbb{R}^m_{>0}$ expresses the relative emphasis on different objectives. Preference-guided optimization seeks relevant compromises among these objectives, rather than requiring every property to improve at each step (Mahapatra & Rajan, 2020). Such preferences guide property-space trade-offs but do not, by themselves, specify which molecular regions should be retained during structural editing.

3.2 DST-BASED MOLECULAR EDITING

Our structural editing builds on the Differentiable Scaffolding Tree (DST) framework (Fu et al., 2022), using the graph-inversion formulation adopted in our implementation (Niu et al., 2025). A molecule is represented at the substructure level using ring and non-ring-atom units. Additional extension positions are attached to leaf units to represent possible growth. The local editing space supports three operations: deleting a leaf unit, replacing it with another vocabulary unit, and attaching a new unit at an extension position.

In our notation, node-type logits $\mathbf{Z}$ and edge logits $\mathbf{A}$ define the relaxed representation

$$\mathbf{X} = \mathrm{softmax}(\mathbf{Z}), \qquad \mathbf{P} = \mathrm{sigmoid}(\mathbf{A}), \tag{3}$$

where the softmax is applied row-wise. A differentiable property predictor provides gradients with respect to editable graph variables while its parameters remain fixed. Internal units and non-editable adjacency entries are held constant during each local update. The optimized representation is decoded into discrete edits; valid candidates are deduplicated and scored by the property evaluators.

This editing mechanism provides the structural search foundation of MolDeriveX. The following section introduces how property-related anchors, soft preservation constraints, and lineage maintenance guide its use in multi-objective optimization.

4 METHODS

Given initial molecules $\mathcal{P}_0$, normalized evaluators $\mathbf{f}$, a fixed differentiable predictor h_ϕ, and an evaluation budget B, MolDeriveX maintains the inheritance state

$$\mathcal{I}_t(G) = \big(G, \mathcal{A}(r), r, o, \mathcal{H}_t(r)\big), \tag{4}$$

where G is the current parent, r its current lineage root, o its original seed, $\mathcal{A}(r)$ the root's property anchors, and $\mathcal{H}_t(r)$ its quality, visitation, and stagnation records. The state separates reusable structural guidance from the changing parent: anchor matches determine retention costs, while lineage records guide parent allocation and rescue. Descendants inherit the lineage reference, while promotion can refresh r and $\mathcal{A}(r)$ without changing o.

Each generation selects parents, matches their anchors, optimizes editable graph variables, and decodes discrete proposals. The predictor h_ϕ guides proposal generation; evaluators $\mathbf{f}$ score discrete molecules for population selection and reporting. Evaluated descendants then update the population and lineage records. The output is the evaluated candidate set and its nondominated subset. Figure 2 and Algorithm 1 summarize this cycle; the following stages specify anchor selection, soft retention, and lineage updates.

4.1 SELECTING PROPERTY-SPECIFIC STRUCTURAL ANCHORS

Each complete seed molecule starts a lineage; the procedure also applies when only one seed is supplied. This search-time set is distinct from predictor-training data and the editing vocabulary. Seeds need not satisfy every objective: anchors identify evidence worth carrying forward from their existing properties (Appendix I).

For a root molecule G, initially a seed or later a promoted descendant, we enumerate candidate substructures $\mathcal{S}(G)$ from rings, fused ring systems, radius-one and radius-two atom environments, and functional-group matches. Candidates must lie within a prescribed size interval. A fragment's isolated score provides local evidence of property support, but its role in the parent also depends on the surrounding atoms and bonds. We assess this context through a deletion-based score:

$$\ell_i(S) = f_i(S), \qquad \Delta_i(S; G) = f_i(G) - f_i(G \setminus S), \tag{5}$$

where $G \setminus S$ denotes the largest connected component remaining after the atoms of S are removed. A positive Δ_i means that the evaluator assigns a lower score after deletion. This is an attribution estimate for the current parent and evaluator, rather than a causal claim about the fragment. If the complement cannot be scored, the local score $\ell_i(S)$ is used as a fallback.

We rank fragments using their positive deletion effect, isolated score, and the parent's property context, with penalties for negative deletion effects and unnecessarily large fragments. The context term combines parent-property scores with physicochemical desirability. Selecting the highest-ranked fragments for each objective gives anchors $\mathcal{A}_i(G)$, which are carried by the lineage and matched to later parents. This turns whole-molecule scores into localized structural guidance. The complete ranking function, its weights, and fallback rules are specified in Appendix B.

4.2 PARETO EDITING WITH SOFT STRUCTURAL RETENTION

The anchors in $\mathcal{I}_t(G)$ must influence the edit itself. Using the relaxed DST representation from Section 3.2, we optimize the node and edge logits $(\mathbf{Z}, \mathbf{A})$ while keeping the property predictor fixed. The anchor masks determine which editable entries incur an additional cost when they move away from the current parent.

The anchors associated with a parent's lineage are matched to the current molecule. Their atom matches are combined into a protected region: any graph unit overlapping a match is marked, as is its extension position when the unit is a leaf. The resulting node mask is $\boldsymbol{m}_V$, and the edge

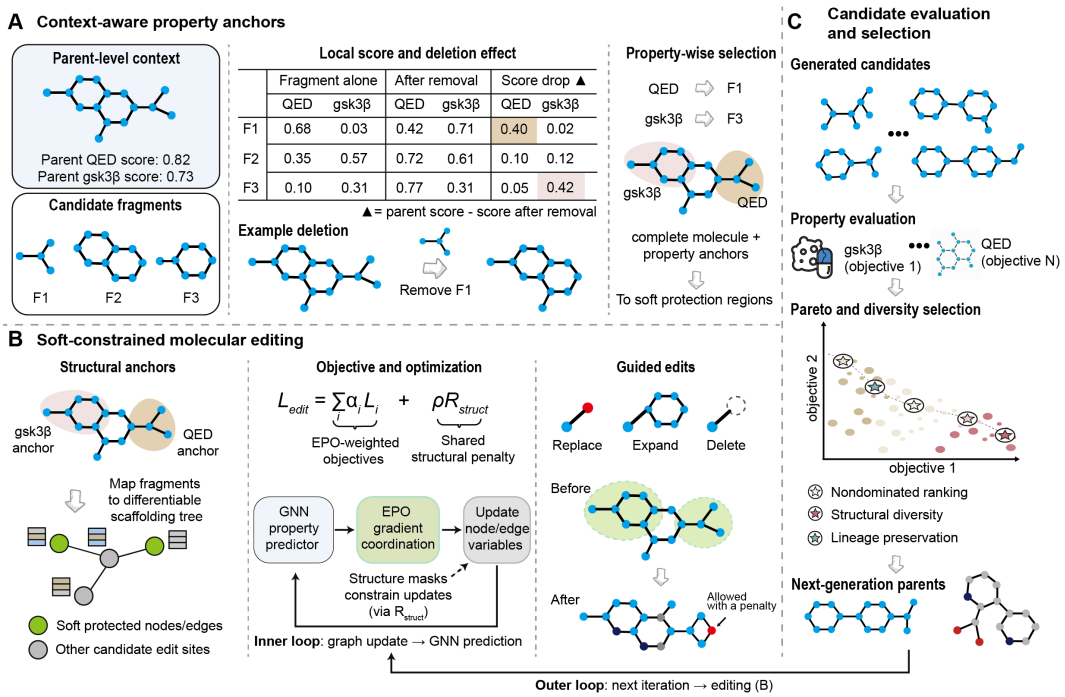


Figure 2: The optimization cycle with adaptive structural inheritance. (A) Anchor selection: fragment scores and parent deletion effects identify property anchors. (B) Soft retention: anchor masks define structural penalties during EPO-coordinated replacement, expansion, and deletion. (C) Continued exploration: evaluated descendants update the population through Pareto rank, structural diversity, and lineage-aware selection, supplying the next generation of parents.

mask $\boldsymbol{M}_E$ marks entries incident to at least one protected node. Anchors are selected separately by property, but their matches share a common structural penalty. The masks are rebuilt for each parent, with its initial logits $(\boldsymbol{Z}^0, \boldsymbol{A}^0)$ serving as the reference:

$$\mathcal{R}_{\text{str}} = \frac{\sum_{v,k} m_{V,v}(Z_{vk} - Z^0_{vk})^2}{\sum_v m_{V,v} + \epsilon} + \frac{\sum_{u,v} M_{E,uv}(A_{uv} - A^0_{uv})^2}{\sum_{u,v} M_{E,uv} + \epsilon}. \tag{6}$$

Here k indexes vocabulary entries and $\epsilon > 0$ stabilizes the normalization. The penalty acts on logits and is normalized by the number of protected nodes and edge entries.

Let $h_{\phi,i}$ be the predictor's output logit for objective i on the relaxed graph. The property loss $\mathcal{L}_i = \log(1 + \exp(-h_{\phi,i}))$ is binary cross-entropy with a target of one. At each edit step, the Exact Pareto Optimal (EPO) search subproblem (Mahapatra & Rajan, 2020) uses these losses, their gradient Gram matrix, and a preference vector $\boldsymbol{r}$ to obtain nonnegative weights $\boldsymbol{\alpha}$ summing to one. We update the graph with

$$\mathcal{L}_{\text{edit}} = m \sum_{i=1}^{m} \alpha_i \mathcal{L}_i + \rho \mathcal{R}_{\text{str}}. \tag{7}$$

If EPO does not return a solution, normalized preference weights are used. The edit representation's fixed internal nodes and fixed adjacency entries are restored after each update.

Equation 7 separates two decisions. The property weights guide the trade-off between objectives, whereas ρ controls the cost of changing the selected structural regions. A finite penalty allows protected, otherwise editable regions to change as the search pursues competing properties. It does not impose a hard scaffold constraint or guarantee simultaneous improvement of every objective.

When adaptive preferences are enabled, objectives below their target scores receive additional weight, blended with a base preference. This directs attention to a parent's weak objectives without replacing Pareto coordination; the exact rule is given in Appendix B.

The optimized graph is decoded through leaf deletion, replacement, and expansion, combining top-k choices with sampling. Actions at protected positions are downweighted by a factor $\kappa \in (0, 1]$; expansion at these positions may be exempted by configuration. Valid proposals are canonicalized and deduplicated. Under budget-capped selection, the predictor ranks new proposals only when they exceed the remaining or per-generation allowance. Evaluators then score the selected proposals; known molecular identities reuse cached scores. These evaluated descendants enter the population update.

4.3 Updating structural references and molecular lineages

A useful local edit should inform subsequent search without making every parent follow the same route. Each initial molecule therefore defines a lineage. We retain both a current root, which can change after promotion, and the original seed identity, which provides a fixed reference for ancestor analysis. When several roots are available, search can follow different structural routes. A single-seed run begins with one original ancestor; descendant promotion can refresh its current roots without adding new initial molecules.

Root prioritization balances incumbent quality with search history. Quality combines mean and minimum objective scores to account for weak objectives. Less-visited roots can receive exploration priority, while stagnation triggers retry, wider edits, weaker retention, or cooldown. These responses let search revisit an unproductive structural reference without requiring another initial molecule.

Separately, when promotion is enabled, descendants meeting configured minimum and mean score thresholds can become new roots, subject to promotion limits. Anchors are recomputed for the promoted molecule, with inherited anchors as fallback; its original seed identity remains unchanged. Promotion thus makes successful modifications available as new structural references, while stagnation rescue addresses unproductive search routes.

Population selection combines retained parents with evaluated descendants. Configured quality quotas, nondominated fronts, crowding, and structural diversity balance objective scores with alternative molecular structures. Limits on parents and survivors per root discourage one route from monopolizing the population and may be relaxed to fill vacant slots. The selected population supplies the next generation; evaluated candidates and their nondominated subset form the output. Search stops at the query or generation limit. Appendix B gives the selection rules and the complete cycle in Algorithm 1.

5 Experiments

The two motivations require evidence of both joint objective quality and productive use of the starting structures. We ask: **(Q1)** how good are the resulting candidate sets; **(Q2)** how do they improve and change structurally relative to their initial ancestors; **(Q3)** how do retention–exploration configurations affect the outcomes; and **(Q4)** can the method obtain useful trade-offs from a single parent in an application case study?

5.1 Experimental Setup

We evaluate six tasks with two to four objectives. Table 1 reports HV on five tasks and Table 2 reports Top-100 quality on the four target activity tasks; complete results are in Appendices F and H. Benchmark searches use 100 initial molecules and a common limit of 5,000 oracle evaluations (Gao et al., 2022). Direct-search baselines comprise MARS (Xie et al., 2021), InversionGNN (Niu et al., 2025), MolSearch (Sun et al., 2022), and GA+DL (Nigam et al., 2020); latent/decoder baselines are MLPS (Liu et al., 2025) and Hgraph (Jin et al., 2020a); policy/exemplar baselines are HN-GFN (Zhu et al., 2023) and RetMol (Wang et al., 2023). Valid molecules are canonicalized, deduplicated, and scored by common frozen evaluators: RDKit QED (Bickerton et al., 2012), normalized SA (Ertl & Schuffenhauer, 2009), normalized logP (RDKit contributors, a), and PyTDC kinase predic-

tors (Huang et al., 2021). Scores are maximized in $[0, 1]$; predicted activity is not an experimental measurement. Model details and transformations appear in Appendices A and D.

We report exact hypervolume (HV) against zero (Zitzler & Thiele, 1999) and the mean of $s(G) = m^{-1} \sum_i f_i(G)$ over the 100 highest-scoring candidates. This Top-100 statistic measures average joint quality, not the fraction of molecules satisfying every objective. In Tables 1 and 2, bold and underlining denote the best and second-best values, ranked before rounding. Additional joint-hit and nonempty Bemis–Murcko scaffold counts (Bemis & Murcko, 1996) use the task-specific thresholds defined in Appendix E.

5.2 Q1: Joint Objective Coverage and Candidate Quality

Table 1: Hypervolume (HV ↑) in five multi-objective molecular optimization tasks.

Method	QED + SA	GSK3β + JNK3	QED + SA + GSK3β	QED + SA + JNK3	QED + SA + GSK3β + JNK3
MARS	0.9438	<u>0.7933</u>	0.8369	**0.8086**	0.5740
InversionGNN	0.7347	0.7361	0.6159	0.4722	0.2777
MolSearch	0.9458	0.6948	0.5688	0.5614	0.3597
GA+DL	0.9456	0.7387	0.6453	0.6480	0.4152
MLPS	0.9375	0.3871	0.5829	0.5385	0.1479
Hgraph	0.9481	0.7850	0.8179	0.6905	<u>0.6281</u>
HN-GFN	0.9392	0.3427	0.4701	0.2597	0.3426
RetMol	<u>0.9481</u>	0.7819	<u>0.8490</u>	0.7355	0.6054
MolDeriveX	**0.9483**	**0.8942**	**0.8779**	<u>0.7694</u>	**0.7103**

MolDeriveX leads HV on four of the five tasks in Table 1, reaching 0.7103 on the four-objective task versus 0.6281 for Hgraph. QED–SA is nearly tied with RetMol, whereas MARS leads QED–SA–JNK3. HV measures frontier coverage; we next examine whether the search also accumulates multiple high-quality candidates.

Table 2: Top-100 joint mean score (↑) on the four target activity tasks.

Method	GSK3β + JNK3	QED + SA + GSK3β	QED + SA + JNK3	QED + SA + GSK3β + JNK3
MARS	0.6631	0.7295	0.6997	0.6290
InversionGNN	0.6440	0.6010	0.5741	0.5339
MolSearch	0.5341	0.5689	0.5367	0.5063
GA+DL	0.6313	0.6584	0.6702	0.5917
MLPS	0.4269	0.6974	0.6498	0.5116
Hgraph	<u>0.7459</u>	0.8056	<u>0.7835</u>	<u>0.7211</u>
HN-GFN	0.4875	0.6187	0.5999	0.4978
RetMol	0.6255	<u>0.8178</u>	0.7265	0.6730
MolDeriveX	**0.7884**	**0.8631**	**0.8066**	**0.8049**

Table 2 evaluates joint candidate quality on tasks including GSK3β, JNK3, or both. MolDeriveX leads all four tasks, reaching 0.8049 on the four-objective task versus 0.7211 for the strongest baseline, Hgraph. Thus, its advantage extends from frontier coverage to the average quality of a set of 100 selected candidates.

On QED–SA–JNK3, MolDeriveX leads Top-100 quality while MARS leads HV. Thus frontier extent and the quality of multiple selected candidates capture different benefits. Top-10 and Top-50 give the same four-task ranking (Appendix H). The two descriptor-only tasks are included there as well: MolDeriveX ranks second and third in Top-100 score, locating its clearest advantage in activity-containing optimization.

High average scores can still mask an unsatisfied objective. On the four-objective task, MolDeriveX also yields 381 joint hits spanning 223 nonempty scaffolds, versus at most 22 hits and 21 scaffolds for a baseline. The joint hits show that high average quality is accompanied by candidates satisfying

every objective, and their scaffolds establish structurally distinct alternatives. Threshold sensitivity and complete candidate diversity, uniqueness, and novelty results appear in Appendices F, G, and M.

5.3 Q2: Improvement beyond the Starting Molecules

Q1 compares final candidate quality across methods; Q2 asks whether MolDeriveX improves on its own starting molecules. We evaluate both objective-space coverage relative to the initial population and candidate scores relative to recorded initial ancestors (Appendix I).

On the four-objective task, the initial 100 molecules have HV 0.3597 and no joint hit. Uniform 100-row subsets of the final output reach mean HV 0.5295, comparing coverage at the same sample size (Appendix K). The complete final set contains 381 joint hits, showing that optimization produces candidates satisfying all four thresholds from a starting set with none.

At the molecule level, every final Top-100 candidate improves its equal-weight mean objective score over its recorded initial ancestor, with an average gain of 0.2837. Of these candidates, 55% also Pareto-dominate their ancestors; the remaining 45% improve the mean score with trade-offs between objectives. Average gains of 0.4091 in GSK3β and 0.4153 in JNK3 show improvement in both activity objectives across this selected set. Together, the set-level and ancestor comparisons establish improvement beyond the supplied starting molecules.

Direct-parent changes and structural-neighborhood comparisons address complementary questions about individual edits and proximity to the starting set (Appendices I.1 and J). Q3 examines objective quality and structural coverage across alternative search configurations.

5.4 Q3: Retention–Exploration Trade-offs

Table 3: Configuration ablation across five tasks (HV ↑).

Configuration	QED + SA	GSK3β + JNK3	QED + SA + GSK3β	QED + SA + JNK3	QED + SA + GSK3β + JNK3
Full	**0.9483**	**0.8942**	**0.8779**	**0.7694**	**0.7103**
Local-only anchoring	0.9418	0.8476	0.7878	0.7440	0.6213
w/o structural retention	0.9461	0.8645	0.8202	0.7348	0.6662
w/o lineage rescue	0.9444	0.7981	0.8505	0.7386	0.6600

Full attains the highest HV on all five tasks (Table 3). Its lead over the strongest alternative is 0.0022 on QED–SA, versus 0.0297 on the dual-kinase task and 0.0441 with four objectives. Local-only anchoring trails Full by 0.0901 on QED–SA–GSK3β and 0.0890 with four objectives, favoring context-informed structural guidance in these comparisons.

Dual-kinase diversity is 0.8481 for Full versus 0.7848 without lineage rescue; four-objective scaffold counts are 3,697 versus 3,292. Full thus combines objective quality with structural exploration without maximizing every diversity or scaffold metric. QED–logP and complete structural results appear in Appendix P.

5.5 Q4: Single-Parent Antibiotic Derivatization Case Study

To examine optimization from a single structural reference, we evaluate three settings initialized from linezolid, clindamycin, or erythromycin, respectively, each with one parent. Objectives are predicted activity from the pretrained Chemprop model of Liu et al. (2023) and Morgan similarity to that parent. The activity endpoint is predicted inhibition of *A. baumannii* (Appendix L). Comparators include neural-guided genetic search (GA+D; Nigam et al., 2020), transformation-based MolSearch (Sun et al., 2022), and F-CReM, the case-specific fragment-editing comparator built on context-dependent CReM mutations (Polishchuk, 2020).

Despite the single-parent initialization, MolDeriveX attains the highest activity–similarity HV on all three tasks: 0.69, 0.69, and 0.79, versus 0.46, 0.54, and 0.70 for the strongest respective comparator (Figure 5). Figure 3 shows higher similarity than GA+D and F-CReM for linezolid; clindamycin favors retention while some competitors reach higher activity at lower similarity, and erythromycin

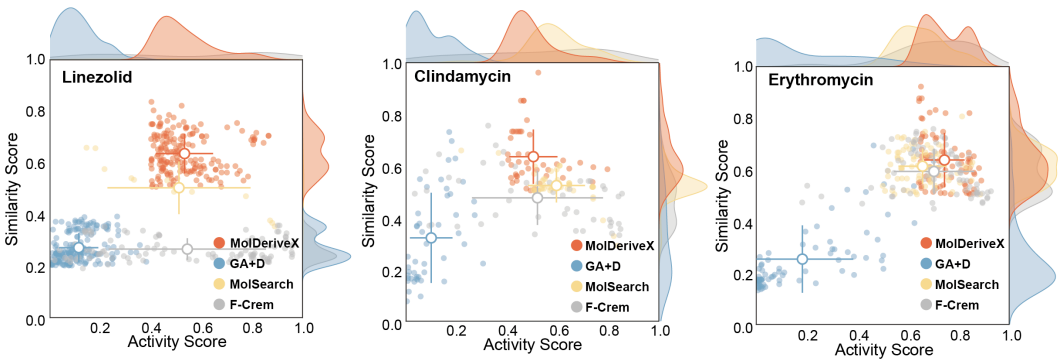


Figure 3: Activity–similarity exploration in three separate single-parent tasks. Points show candidate scores; marginal curves show distributions. Open circles and horizontal/vertical bars denote candidate-score means and standard deviations.

shows greater method overlap. The advantage is joint trade-off quality. Diversity of extracted derivative regions is also highest for MolDeriveX in these comparisons, indicating varied modifications from one parent (Appendix L). This metric excludes the fixed core and differs from whole-molecule diversity.

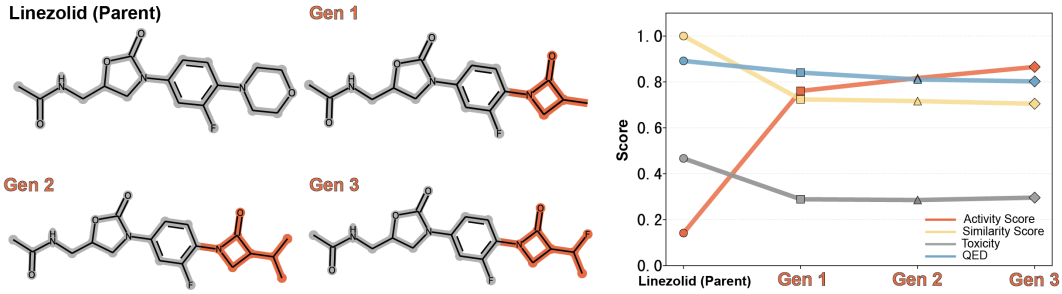


Figure 4: Linezolid derivatization sequence. Gray and orange mark shared and modified regions. The right panel shows predicted antibacterial activity, parent similarity, predicted toxicity, and QED.

Figure 4 illustrates structural continuity and change along one derivatization sequence. Predicted activity rises from approximately 0.14 to 0.86 as similarity falls from 1.00 to about 0.70. The largest activity increase accompanies the first revision; later descendants retain a common region while extending the modification. QED decreases modestly, and the predicted toxicity score remains below the parent's. This sequence shows how improved predicted activity can coexist with structural continuity and changes in other properties. The trajectory complements the candidate-set comparisons by showing how successive edits retain shared structure while improving predicted activity.

6 Conclusion

MolDeriveX uses adaptive structural inheritance to connect property-specific anchors, soft Pareto editing, and lineage-reference updates. Benchmark quality and ancestor gains show productive optimization from limited starting structures, while configuration comparisons reveal trade-offs between objective quality and structural coverage. Three single-parent antibiotic tasks extend this evidence to predicted activity–similarity optimization. Together, these results support reusing structural information while allowing its revision as objectives compete. Population-size sweeps, isolated component tests, and experimental activity measurements are needed to establish initialization sensitivity, causal contributions, and biological efficacy, respectively.

7 AI USE STATEMENT

During code development, AI tools were used for debugging and refinement of the experimental code. For scientific figures, generative AI tools were used to create or refine selected visual elements. AI tools were also used for language editing to improve clarity and readability. All AI-assisted outputs were reviewed and verified by the authors before being incorporated into this work.

REFERENCES

Guy W. Bemis and Mark A. Murcko. The properties of known drugs. 1. molecular frameworks. *Journal of Medicinal Chemistry*, 39(15):2887–2893, 1996. doi: 10.1021/jm9602928.

G Richard Bickerton, Gaia V Paolini, Jérémy Besnard, Sorel Muresan, and Andrew L Hopkins. Quantifying the chemical beauty of drugs. *Nature Chemistry*, 4(2):90–98, 2012. doi: 10.1038/nchem.1243.

Yair Censor. Pareto optimality in multiobjective problems. *Applied Mathematics and Optimization*, 4(1):41–59, 1977. doi: 10.1007/BF01442131.

Peter Ertl and Ansgar Schuffenhauer. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *Journal of Cheminformatics*, 1:8, 2009. doi: 10.1186/1758-2946-1-8.

Tianfan Fu, Cao Xiao, Xinhao Li, Lucas M Glass, and Jimeng Sun. MIMOSA: Multi-constraint molecule sampling for molecule optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 125–133, 2021. doi: 10.1609/aaai.v35i1.16085.

Tianfan Fu, Wenhao Gao, Cao Xiao, Jacob Yasonik, Connor W. Coley, and Jimeng Sun. Differentiable Scaffolding Tree for molecule optimization. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=w_drCosT76.

Wenhao Gao, Tianfan Fu, Jimeng Sun, and Connor Coley. Sample efficiency matters: a benchmark for practical molecular optimization. *Advances in Neural Information Processing Systems*, 35: 21342–21357, 2022. URL https://arxiv.org/abs/2206.12411.

Rafael Gómez-Bombarelli, Jennifer N Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D Hirzel, Ryan P Adams, and Alán Aspuru-Guzik. Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Science*, 4(2):268–276, 2018. doi: 10.1021/acscentsci.7b00572.

Ryan-Rhys Griffiths and José Miguel Hernández-Lobato. Constrained bayesian optimization for automatic chemical design using variational autoencoders. *Chemical Science*, 11(2):577–586, 2020. doi: 10.1039/C9SC04026A.

Kexin Huang, Tianfan Fu, Wenhao Gao, Yue Zhao, Yusuf Roohani, Jure Leskovec, Connor W. Coley, Cao Xiao, Jimeng Sun, and Marinka Zitnik. Therapeutics Data Commons: Machine learning datasets and tasks for drug discovery and development. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1, 2021. URL https://arxiv.org/abs/2102.09548.

Jan H Jensen. A graph-based genetic algorithm and generative model/monte carlo tree search for the exploration of chemical space. *Chemical Science*, 10(12):3567–3572, 2019. doi: 10.1039/C8SC05372C.

Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Junction tree variational autoencoder for molecular graph generation. In *International conference on machine learning*, pp. 2323–2332. PMLR, 2018. URL https://proceedings.mlr.press/v80/jin18a.html.

Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Hierarchical generation of molecular graphs using structural motifs. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pp. 4839–4848. PMLR, 2020a. URL https://proceedings.mlr.press/v119/jin20a.html.

Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Multi-objective molecule generation using interpretable substructures. In *International conference on machine learning*, pp. 4849–4859. PMLR, 2020b. URL https://proceedings.mlr.press/v119/jin20b.html.

Gary Liu, Denise B. Catacutan, Khushi Rathod, Kyle Swanson, Wengong Jin, Jody C. Mohammed, Anush Chiappino-Pepe, Saad A. Syed, Meghan Fragis, Kenneth Rachwalski, Jakob Magolan, Michael G. Surette, Brian K. Coombes, Tommi Jaakkola, Regina Barzilay, James J. Collins, and Jonathan M. Stokes. Deep learning-guided discovery of an antibiotic targeting Acinetobacter baumannii. *Nature Chemical Biology*, 19:1342–1350, 2023. doi: 10.1038/s41589-023-01349-8.

Yiping Liu, Jiahao Yang, Xuanbai Ren, Zhang Xinyi, Yuansheng Liu, Bosheng Song, Xiangxiang Zeng, and Hisao Ishibuchi. Multi-objective molecular design through learning latent pareto set. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 19006–19014, 2025. doi: 10.1609/aaai.v39i18.34092.

Debabrata Mahapatra and Vaibhav Rajan. Multi-task learning with user preferences: Gradient descent with controlled ascent in Pareto optimization. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pp. 6597–6607. PMLR, 2020. URL https://proceedings.mlr.press/v119/mahapatra20a.html.

Gergely M. Makara, László Kovács, István Szabó, and Gábor Pőcze. Derivatization design of synthetically accessible space for optimization: In silico synthesis vs deep generative design. *ACS Medicinal Chemistry Letters*, 12(2):185–194, 2021. doi: 10.1021/acsmedchemlett.0c00540.

AkshatKumar Nigam, Pascal Friederich, Mario Krenn, and Alán Aspuru-Guzik. Augmenting genetic algorithms with deep neural networks for exploring the chemical space. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=H1lmyRNFvr.

Yifan Niu, Ziqi Gao, Tingyang Xu, Yang Liu, Yatao Bian, Yu Rong, Junzhou Huang, and Jia Li. InversionGNN: A dual path network for multi-property molecular optimization. In *International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=nYPuSzGE3X.

Pavel Polishchuk. CReM: chemically reasonable mutations framework for structure generation. *Journal of Cheminformatics*, 12:28, 2020. doi: 10.1186/s13321-020-00431-w.

RDKit contributors. RDKit documentation: rdkit.Chem.Crippen. https://www.rdkit.org/docs/source/rdkit.Chem.Crippen.html, a. Accessed September 25, 2026.

RDKit contributors. RDKit documentation: rdkit.Chem.rdfiltercatalog. https://www.rdkit.org/docs/source/rdkit.Chem.rdfiltercatalog.html, b. Accessed September 25, 2026.

David Rogers and Mathew Hahn. Extended-connectivity fingerprints. *Journal of Chemical Information and Modeling*, 50(5):742–754, 2010. doi: 10.1021/ci100050t.

Mengying Sun, Jing Xing, Han Meng, Huijun Wang, Bin Chen, and Jiayu Zhou. MolSearch: Search-based multi-objective molecular generation and property optimization. In *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*, pp. 4724–4732, 2022. doi: 10.1145/3534678.3542676.

Therapeutics Data Commons. Oracle functions. https://tdcommons.ai/functions/oracles/. Accessed September 25, 2026.

Paula Torren-Peraire, Nadine Schneider, Raquel Rodríguez-Pérez, Robin A Fairhurst, and Jessica Lanini. Multi-objective optimization in the context of generative chemistry. *Nature Communications*, 17(1):9405, 2026.

Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.

Jonas Verhellen. Graph-based molecular pareto optimisation. *Chemical Science*, 13(25):7526–7535, 2022. doi: 10.1039/D2SC00821A.

Zichao Wang, Weili Nie, Zhuoran Qiao, Chaowei Xiao, Richard G. Baraniuk, and Anima Anandkumar. Retrieval-based controllable molecule generation. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=vDFA1tpuLvk.

Zhenqin Wu, Bharath Ramsundar, Evan N. Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S. Pappu, Karl Leswing, and Vijay Pande. MoleculeNet: a benchmark for molecular machine learning. *Chemical Science*, 9:513–530, 2018. doi: 10.1039/C7SC02664A.

Yutong Xie, Chence Shi, Hao Zhou, Yuwei Yang, Weinan Zhang, Yong Yu, and Lei Li. Mars: Markov molecular sampling for multi-objective drug discovery. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=kHSu4ebxFXY.

Kevin Yang, Kyle Swanson, Wengong Jin, Connor Coley, Philipp Eiden, Hua Gao, Angel Guzman-Perez, Timothy Hopper, Brian Kelley, Miriam Mathea, Andrew Palmer, Volker Settels, Tommi Jaakkola, Klavs Jensen, and Regina Barzilay. Analyzing learned molecular representations for property prediction. *Journal of Chemical Information and Modeling*, 59(8):3370–3388, 2019. doi: 10.1021/acs.jcim.9b00237.

Jiaxuan You, Bowen Liu, Zhitao Ying, Vijay Pande, and Jure Leskovec. Graph convolutional policy network for goal-directed molecular graph generation. *Advances in neural information processing systems*, 31, 2018. URL https://arxiv.org/abs/1806.02473.

Xiangxiang Zeng, Fei Wang, Yuan Luo, Seung-gu Kang, Jian Tang, Felice C Lightstone, Evandro F Fang, Wendy Cornell, Ruth Nussinov, and Feixiong Cheng. Deep generative molecular design reshapes drug discovery. *Cell Reports Medicine*, 3(12), 2022.

Zhenpeng Zhou, Steven Kearnes, Li Li, Richard N Zare, and Patrick Riley. Optimization of molecules via deep reinforcement learning. *Scientific Reports*, 9(1):10752, 2019. doi: 10.1038/s41598-019-47148-x.

Yiheng Zhu, Jialu Wu, Chaowen Hu, Jiahuan Yan, Kim Hsieh, Tingjun Hou, and Jian Wu. Sample-efficient multi-objective molecular optimization with GFlowNets. *Advances in Neural Information Processing Systems*, 36:79667–79684, 2023. doi: 10.52202/075280-3487.

Eckart Zitzler and Lothar Thiele. Multiobjective evolutionary algorithms: A comparative case study and the strength Pareto approach. *IEEE Transactions on Evolutionary Computation*, 3(4):257–271, 1999. doi: 10.1109/4235.797969.

A DIFFERENTIABLE PROPERTY PREDICTOR

MolDeriveX uses a differentiable graph predictor to guide structural edits and separate property oracles to evaluate discrete molecules. The predictor supplies gradients and proposal-ranking scores; the benchmark results use the property evaluators defined in Appendix D.

The predictor takes a substructure-vocabulary matrix X, adjacency matrix A, and node-weight vector w. A linear layer embeds vocabulary coordinates into 50 dimensions. An input graph convolution maps this embedding to 100 hidden dimensions, followed by three additional 100-dimensional graph convolutions with ReLU activation. An output graph convolution produces one logit per modeled property. Each convolution computes $AHW + b$. The graph-level output is a node-weighted mean:

$$h_{\phi,i}(G) = \frac{\sum_v w_v H_{vi}^{\text{out}}}{\sum_v w_v}. \tag{8}$$

Sigmoid maps the logits to proposal scores in $[0, 1]$. The network operates on ring and non-ring-atom units, following the differentiable graph-editing representation (Fu et al., 2022; Niu et al., 2025). During molecular editing, network parameters remain fixed and gradients update the editable graph logits. The EPO objective uses pre-sigmoid property outputs, with output coordinates aligned to the task objectives.

B IMPLEMENTATION DETAILS OF STRUCTURAL GUIDANCE

Fragment context and ranking. We also compute a context score $c_i(G) \in [0, 1]$ from the parent-property score, QED, and smooth desirability functions of lipophilicity and heavy-atom count. Candidate fragments are ranked by

$$\begin{aligned} R_i(S; G) = w_I \left[\Delta_i(S; G)(0.5 + 0.5c_i(G))\right]_+ + w_L \ell_i(S) + w_G c_i(G) \\ - w_C[-\Delta_i(S; G)]_+ - w_S \sqrt{\min(1, |S|/s_{\max})}, \end{aligned} \tag{9}$$

where $[x]_+ = \max(0, x)$ and the nonnegative coefficients weight positive contextual contribution, local support, parent context, negative contextual contribution, and fragment size. The size term discourages retaining unnecessarily large regions. For a fixed parent and objective, the additive context term is constant across candidates; context affects their relative ranking through its interaction with Δ_i. The highest-ranked fragments are retained for each objective, producing a set $\mathcal{A}_i(G)$ of structural anchors. An optional local-score fallback favours the best local candidate when the contextual candidate does not exceed it by a specified margin under the same retention score.

For objective i, the context score is

$$c_i(G) = \text{clip}\left(0.55 f_i(G) + 0.20\,\text{QED}(G) + 0.15 h_{\text{logP}}(G) + 0.10 h_{\text{size}}(G), 0, 1\right), \tag{10}$$

where $h_{\text{logP}} = \exp[-\frac{1}{2}((\text{logP} - 2.5)/2.5)^2]$ and $h_{\text{size}} = \exp[-\frac{1}{2}((n_{\text{heavy}} - 30)/22)^2]$. This context desirability differs from the monotone clipped logP reward used in the QED–logP benchmark. The default ranking weights in Equation 9 are $(w_I, w_L, w_G, w_C, w_S) = (0.55, 0.20, 0.20, 0.50, 0.05)$. The default extraction settings consider fragments with 3–18 atoms and retain one hit per objective. Fragment deletion retains the largest scorable complement; a local-score fallback is used when that complement cannot be evaluated. These values specify the default anchor-extraction rule; the ablation contrasts are described in Appendix P.

Adaptive objective preferences. When adaptive preferences are enabled, a parent's deficit $d_i(G) = [\tau_i - f_i(G)]_+$ determines an additional emphasis on objective i:

$$\begin{aligned} u_i(G) &= \text{clip}\big(d_i(G)^\gamma + r_{\min}, r_{\min}, r_{\max}\big), \\ \boldsymbol{r}(G) &= (1 - \beta)\boldsymbol{r}_0 + \beta \frac{\boldsymbol{u}(G)}{\sum_j u_j(G)}. \end{aligned} \tag{11}$$

The normalized base preference $\boldsymbol{r}_0$ is retained if all deficits vanish. Population-level deficit feedback can update the base search preference between generations.

Optimization procedure. Algorithm 1 summarizes the initialization, local editing, oracle evaluation, and population update. Repeated molecular identities reuse cached scores. The remaining query allowance limits only new oracle evaluations; preprocessing and post-hoc analysis are separate from the search loop. With strict budget control, predictor-based proposal ranking is invoked only if new identities exceed the remaining or per-generation query allowance; otherwise all new proposals proceed to evaluation.

Algorithm 1 MolDeriveX: soft-constrained editing with lineage exploration

Require: Initial population $\mathcal{P}_0$, predictor h_ϕ, oracles $\mathbf{f}$, remaining query allowance B, generation limit T
Ensure: Evaluated proposals $\mathcal{E}$ and their nondominated subset

Extract property anchors and initialize lineage roots for $\mathcal{P}_0$
Initialize score cache $\mathcal{C}$ with available seed scores
$\mathcal{E} \leftarrow \emptyset$; $b \leftarrow B$; $t \leftarrow 0$
while $b > 0$ and $t < T$ **do**
    Select parents using population and root-priority rules
    $\mathcal{V} \leftarrow \emptyset$
    for each selected parent G **do**
        Match lineage anchors to G; construct protected graph masks
        Initialize editable logits from G and choose property preferences
        for each local update step **do**
            Obtain EPO weights; update logits using Equation 7
            Restore fixed graph entries
        end for
        Decode deletion, replacement, and expansion proposals into $\mathcal{V}$
    end for
    Canonicalize valid proposals and remove duplicates from $\mathcal{V}$
    if $\mathcal{V} = \emptyset$ **then**
        break
    end if
    $\mathcal{N} \leftarrow \{G \in \mathcal{V} : G \notin \mathrm{dom}(\mathcal{C})\}$
    Set a to b capped by any per-generation allowance; $\mathcal{Q} \leftarrow \mathcal{N}$
    if $|\mathcal{N}| > a$ **then**
        Select $\mathcal{Q} \subseteq \mathcal{N}$, $|\mathcal{Q}| \leq a$, using predictor scores and allocation rules
    end if
    Evaluate $\mathbf{f}(G)$ for $G \in \mathcal{Q}$; extend $\mathcal{C}$; $b \leftarrow b - |\mathcal{Q}|$
    $\mathcal{V}_{\mathrm{eval}} \leftarrow \mathcal{V} \cap \mathrm{dom}(\mathcal{C})$; $\mathcal{E} \leftarrow \mathcal{E} \cup \mathcal{V}_{\mathrm{eval}}$
    Update root incumbents; apply enabled stagnation responses
    Promote eligible descendants and refresh their anchors if enabled
    Select $\mathcal{P}_{t+1}$ from $\mathcal{P}_t \cup \mathcal{V}_{\mathrm{eval}}$ using quality, diversity, and root limits
    $t \leftarrow t + 1$
end while
return $\mathcal{E}$ and $\mathrm{ND}(\mathcal{E})$

Root quality and allocation. For a molecule G, define balanced quality $q_\omega(G) = (1 - \omega)m^{-1}\sum_i f_i(G) + \omega \min_i f_i(G)$. The minimum-score term discourages an excellent score on one objective from masking a weak objective. When root prioritization is enabled, the implementation assigns a root r priority

$$U(r) = q_r + 0.25b_r + c\sqrt{\frac{\log(N_t + 2)}{n_r + 1}} - \zeta s_r + \xi_r, \tag{12}$$

where q_r and b_r track incumbent balanced quality and minimum score, n_r is the root's visit count, N_t sums visits across active roots, and s_r counts stagnation. The configured exploration coefficient c balances incumbent quality with less-visited roots; ξ_r collects optional retry, promotion, and cooldown adjustments.

Promotion and population selection. Stagnation controls retry, retention relaxation, and cooldown. Separately, when root promotion is enabled, candidates meeting configured minimum objective-score and mean-score thresholds are ranked by those two scores and promoted subject to per-generation and total limits. The promoted molecule becomes a new root and receives newly extracted anchors, or inherits its previous root's anchors if extraction yields no replacement. Its original ancestor identity is preserved throughout promotion.

The Pareto-diverse selection mode pools retained parents and evaluated descendants. Configured quotas for original-ancestor coverage, balanced scores, and deficient objectives are filled first; remaining slots are filled by nondominated fronts. Within a front, the score combines structural diversity, estimated HV contribution, crowding, mean score, and minimum score, with configured weights. Structural diversity relative to the selected set $\mathcal{S}$ is $1 - \max_{H \in \mathcal{S}} \mathrm{Tan}(G, H)$, with value one for an empty set. Per-root limits apply during selection and may be relaxed to fill vacant slots. Each configuration specifies its active quotas and weighting terms.

Structural guidance during editing. The retention penalty acts on graph logits, while protected-position downweighting acts during decoding. Editable anchor regions can change or be removed as the search balances the objectives. Promotion can establish a new root with refreshed anchors while preserving the molecule's original seed identity. Accordingly, nearest-initial fingerprint similarity, direct-parent changes, and initial-ancestor changes answer different questions and are reported separately in the supplementary analyses.

C Evaluation tasks and protocol

We compare nine methods on six combinations of molecular properties (Table 4), with a search budget of 5,000 oracle evaluations and 100 initial molecules. The evaluation contains one candidate set per method and task. All candidate sets use the same property evaluators, normalization, molecular-identity rules, and success thresholds.

Table 4: Task definitions and objective order. All normalized objectives are maximized.

Task	Ordered objectives
QED–logP	QED, logP reward
QED–SA	QED, SA reward
GSK3B–JNK3	GSK3B, JNK3
QED–SA–GSK3B	QED, SA reward, GSK3B
QED–SA–JNK3	QED, SA reward, JNK3
QED–SA–GSK3B–JNK3	QED, SA reward, GSK3B, JNK3

The comparison includes MARS, InversionGNN, MolSearch, GA+DL, MLPS, Hgraph, HN-GFN, RetMol, and MolDeriveX. In the ablation tables, Full denotes the complete MolDeriveX configuration.

C.1 Molecular identity and eligibility

SMILES are parsed and sanitized with RDKit and converted to canonical isomeric SMILES. Charges, stereochemistry, and all disconnected fragments are retained. Empty or unparsable strings are invalid. Molecular deduplication is performed separately within each method–task collection. Different molecular identities with identical score vectors remain distinct for hit counts and structural analyses.

Let R be a collection of raw records, $N = |R|$, and V its number of valid records, including duplicates. Let U denote the set of valid, unique molecules with complete task scores. All valid molecules in this evaluation have complete task scores. Failed scores are marked as unavailable. Raw-record denominators and output subsampling include invalid records; final-set metrics are computed from the exported candidates.

D PROPERTY EVALUATORS AND SCORE TRANSFORMATIONS

We evaluate QED, synthetic accessibility, lipophilicity, and predicted GSK3β/JNK3 activity. The same scoring functions and transformations are applied to all methods. The property predictor used for graph editing is separate from these evaluators.

D.1 QED

For molecular graph x, the QED objective is

$$f_{\mathrm{QED}}(x) = \texttt{rdkit.Chem.QED.qed}(x). \tag{13}$$

The standard RDKit QED value is used directly, without fitting an additional normalization to any output set. QED summarizes selected molecular descriptors through the quantitative estimate of drug-likeness Bickerton et al. (2012).

D.2 SYNTHETIC ACCESSIBILITY REWARD

The raw synthetic-accessibility score $a(x)$ uses the RDKit implementation of the fragment-contribution and molecular-complexity measure of Ertl & Schuffenhauer (2009). The maximization reward is

$$f_{\mathrm{SA}}(x) = \exp\left[-\frac{1}{2}\left(\frac{\max\{a(x),\mu\}-\mu}{\sigma}\right)^2\right], \qquad \mu = 2.230044, \quad \sigma = 0.6526308. \tag{14}$$

These constants are fixed across methods and tasks. Lower raw SA values indicate greater estimated accessibility, whereas the transformed reward is maximized. Every raw value $a(x) \le \mu$ maps to a reward of 1.

D.3 LIPOPHILICITY REWARD

Let $\ell(x)$ be the RDKit Wildman–Crippen logP estimate returned by `Crippen.MolLogP` RDKit contributors (a). The objective is

$$f_{\mathrm{logP}}(x) = \min\left\{1, \max\left\{0, \frac{\ell(x)+10}{14}\right\}\right\}. \tag{15}$$

It saturates at 0 for $\ell \le -10$ and at 1 for $\ell \ge 4$. The transformation defines a monotone clipped reward. Raw logP is retained separately for descriptive quality analysis.

D.4 GSK3B AND JNK3 ACTIVITY SCORES

Following multi-objective molecular optimization benchmarks (Jin et al., 2020b; Wang et al., 2023), we evaluate predicted GSK3β and JNK3 activity using the pretrained `GSK3B` and `JNK3` oracles provided by PyTDC (Huang et al., 2021; Therapeutics Data Commons). Their scores lie in $[0, 1]$, with larger values preferred. We use the same pretrained oracles for every method.

E METRIC DEFINITIONS

E.1 JOINT SUCCESS AND DENOMINATORS

For task objective set J and threshold vector $\boldsymbol{\tau}$, define

$$h(x;\boldsymbol{\tau}) = \prod_{j\in J} \mathbf{1}\{f_j(x) \ge \tau_j\}, \qquad H = \{x \in U : h(x;\boldsymbol{\tau}) = 1\}. \tag{16}$$

Every objective must pass; the threshold operator is non-strict. The three predefined profiles are listed in Table 5. Thresholds apply to normalized objective scores.

Table 5: Joint-success thresholds. Only the objectives included in a task are applied. All comparisons use $\geq$.

Profile	QED	SA reward	logP reward	GSK3B	JNK3
Main	0.60	0.67	0.80	0.50	0.50
Moderate	0.70	0.80	0.85	0.60	0.60
Strict	0.80	0.90	0.90	0.70	0.70

We report the number of distinct valid molecules satisfying every threshold, $|H|$.

E.2 Hypervolume and Pareto coverage

All objectives have a fixed $[0, 1]$ scale and are maximized. Standard hypervolume is the volume of the union of origin-anchored boxes:

$$\mathrm{HV}(U) = \mathrm{vol}_{|J|}\left(\bigcup_{x\in U}\prod_{j\in J}[0, f_j(x)]\right). \tag{17}$$

HV uses all valid unique score-complete output molecules, with empty-set HV defined as zero. Duplicate score vectors and zero-volume boxes are removed before exact computation with PyGMO: `hv2d` for two objectives, `hv3d` for three, and `hvwfg` for four. Method comparisons use the same task objectives and reference point.

E.3 Fingerprint and scaffold diversity

For a molecular set S of size $n \geq 2$, internal diversity is

$$D(S) = 1 - \frac{2}{n(n-1)}\sum_{i<j}\mathrm{Tan}(\phi(x_i), \phi(x_j)). \tag{18}$$

Structural comparisons use 2,048-bit Morgan fingerprints with radius 2, bond types enabled, and chirality disabled. Every unordered pair of distinct molecules is included exactly, with self-pairs excluded. Diversity is undefined for $n < 2$. The reported overall internal diversity uses the set of all valid unique molecules.

Scaffolds are non-isomeric canonical RDKit Bemis–Murcko scaffolds. Successful scaffold coverage counts distinct nonempty scaffolds among joint hits. It measures absolute coverage and is affected by the number of successful molecules.

E.4 Reporting conventions

Tables summarize the evaluated candidate sets. A dash denotes an unavailable value. Output-subset ranges follow the resampling procedure in Appendix K; all rankings use unrounded scores. Bold denotes the best value and underlining the second distinct value.

F Objective coverage and successful structural coverage

Table 1 and the descriptor-task summary in Table 10 cover all six HV comparisons. MolDeriveX leads on GSK3B–JNK3, QED–SA, QED–SA–GSK3B, and the four-objective task. MARS leads QED–SA–JNK3, while MolDeriveX is third on QED–logP. The QED–SA margin over the strongest baseline is only 1.21×10^{-4}.

The four target activity tasks show the clearest advantage in successful candidate production (Table 6, reported as unique joint hits / distinct nonempty hit scaffolds). On the four-objective task, 381 hits span 223 nonempty scaffolds, compared with at most 22 hits and 21 scaffolds for a baseline. Thus the high joint scores extend to multiple successful structures. Overall fingerprint diversity is 0.7897 for MolDeriveX versus 0.8738 for Hgraph, highlighting the distinction between successful

scaffold coverage and global structural dispersion. Descriptor-only results remain visible in Table 10.

Table 6: Joint hits and successful scaffold coverage: target activity tasks.

Method	GSK3β + JNK3	QED + SA + GSK3β	QED + SA + JNK3	QED + SA + GSK3β + JNK3
MARS	90 / 66	27 / 22	14 / 10	3 / 3
InversionGNN	222 / 153	0 / 0	0 / 0	0 / 0
MolSearch	2 / 2	0 / 0	0 / 0	0 / 0
GA+DL	150 / 84	0 / 0	0 / 0	0 / 0
MLPS	0 / 0	9 / 8	0 / 0	0 / 0
Hgraph	224 / 106	72 / 47	80 / 32	22 / 21
HN-GFN	1 / 1	0 / 0	0 / 0	1 / 1
RetMol	73 / 58	129 / 78	16 / 10	9 / 9
MolDeriveX	1217 / 695	745 / 287	315 / 92	381 / 223

G SENSITIVITY TO JOINT-SUCCESS THRESHOLDS

Tables 7 and 8 report unique-hit counts in main / moderate / strict profile order for every method. Tightening the four-objective thresholds from the main to the moderate profile leaves 110 MolDeriveX hits versus at most 6 for a baseline. The strict profile leaves only three four-objective hits and none on QED–SA–JNK3. The corresponding dual-kinase and QED–SA–GSK3B counts are 136 and 17. These profiles reveal both the persistence of the advantage and the remaining difficulty of satisfying stricter joint requirements.

Table 7: Sensitivity of joint hits to threshold stringency: target activity tasks.

Method	GSK3β + JNK3	QED + SA + GSK3β	QED + SA + JNK3	QED + SA + GSK3β + JNK3
MARS	90 / 6 / 0	27 / 8 / 0	14 / 6 / 0	3 / 0 / 0
InversionGNN	222 / 2 / 0	0 / 0 / 0	0 / 0 / 0	0 / 0 / 0
MolSearch	2 / 0 / 0	0 / 0 / 0	0 / 0 / 0	0 / 0 / 0
GA+DL	150 / 0 / 0	0 / 0 / 0	0 / 0 / 0	0 / 0 / 0
MLPS	0 / 0 / 0	9 / 0 / 0	0 / 0 / 0	0 / 0 / 0
Hgraph	224 / 122 / 53	72 / 36 / 2	80 / 32 / 1	22 / 6 / 0
HN-GFN	1 / 0 / 0	0 / 0 / 0	0 / 0 / 0	1 / 0 / 0
RetMol	73 / 16 / 2	129 / 62 / 7	16 / 4 / 0	9 / 2 / 0
MolDeriveX	1217 / 422 / 136	745 / 311 / 17	315 / 90 / 0	381 / 110 / 3

Table 8: Sensitivity of joint hits to threshold stringency: descriptor-only tasks.

Method	QED + logP	QED + SA
MARS	365 / 164 / 54	564 / 312 / 115
InversionGNN	371 / 8 / 0	0 / 0 / 0
MolSearch	4 / 1 / 0	172 / 147 / 109
GA+DL	616 / 207 / 25	530 / 315 / 117
MLPS	960 / 647 / 308	327 / 205 / 68
Hgraph	3368 / 2662 / 1560	2438 / 1986 / 1392
HN-GFN	627 / 275 / 74	565 / 222 / 50
RetMol	2804 / 1914 / 856	3126 / 2473 / 1672
MolDeriveX	2747 / 2048 / 1126	728 / 530 / 298

H ELITE CANDIDATE QUALITY

For $d = |J|$, define the equal-weight joint score $\bar{f}(x) = d^{-1} \sum_{j \in J} f_j(x)$. Candidates are ranked by $\bar{f}$ and ties are resolved by ascending canonical SMILES. The reported statistic is

$$Q_k = \frac{1}{k} \sum_{x \in \mathrm{Top}_k(\bar{f};U)} \bar{f}(x), \qquad k \in \{10, 50, 100\}. \tag{19}$$

All objectives are evaluated on the same selected molecules, using equal weights in the reporting score. Top-k is reported for collections with at least k eligible candidates.

The elite-set table reports Top-10 / Top-50 joint mean scores; Top-100 results appear in Table 2. MolDeriveX has the highest joint Top-10, Top-50, and Top-100 means among the evaluated methods on all four target activity tasks. On QED–SA–GSK3B–JNK3, the Top-10 mean is 0.836028 and the Top-100 mean is 0.804855. On QED–logP, MolDeriveX ranks second in Top-100 score (0.957999), behind Hgraph (0.959211; Table 10). On QED–SA, it ranks third (0.960627), behind Hgraph (0.971687) and RetMol (0.971202; Table 10). The largest gains occur on the target activity tasks.

Table 9: Elite-set sensitivity on target activity tasks.

Method	GSK3β + JNK3	QED + SA + GSK3β	QED + SA + JNK3	QED + SA + GSK3β + JNK3
MARS	0.7080 / 0.6791	0.8460 / 0.7714	0.8216 / 0.7379	0.7255 / 0.6603
InversionGNN	0.6845 / 0.6574	0.6784 / 0.6325	0.6536 / 0.6068	0.6152 / 0.5619
MolSearch	0.6680 / 0.5953	0.6718 / 0.6357	0.6648 / 0.6184	0.6482 / 0.5800
GA+DL	0.6785 / 0.6567	0.6869 / 0.6672	0.7205 / 0.6872	0.6482 / 0.6120
MLPS	0.5640 / 0.4726	0.7578 / 0.7245	0.6952 / 0.6654	0.5874 / 0.5384
Hgraph	0.7930 / 0.7642	0.8717 / 0.8408	0.8317 / 0.8097	0.8055 / 0.7561
HN-GFN	0.5285 / 0.5024	0.6845 / 0.6420	0.6357 / 0.6114	0.5852 / 0.5213
RetMol	0.7195 / 0.6625	0.8816 / 0.8405	0.8006 / 0.7542	0.7502 / 0.6963
MolDeriveX	**0.8465 / 0.8135**	**0.8905 / 0.8734**	**0.8532 / 0.8238**	**0.8360 / 0.8177**

Table 10: Complementary results on descriptor-only tasks.

Method	QED–logP HV	QED–logP Top-100	QED–logP Hits / scaffolds	QED–SA HV	QED–SA Top-100	QED–SA Hits / scaffolds
MARS	0.9280	0.8886	365 / 223	0.9438	0.9247	564 / 315
InversionGNN	0.7498	0.7866	371 / 62	0.7347	0.6719	0 / 0
MolSearch	0.6611	0.4411	4 / 2	0.9458	0.9444	172 / 97
GA+DL	0.9061	0.8675	616 / 540	0.9456	0.9223	530 / 268
MLPS	0.9388	0.9464	960 / 678	0.9375	0.9081	327 / 117
Hgraph	0.9476	**0.9592**	3368 / 2185	0.9481	**0.9717**	2438 / 1157
HN-GFN	0.9426	0.9092	627 / 481	0.9392	0.9020	565 / 369
RetMol	**0.9476**	0.9496	2804 / 2077	0.9481	0.9712	3126 / 1848
MolDeriveX	0.9470	0.9580	2747 / 2194	**0.9483**	0.9606	728 / 403

I INITIALIZATION AND IMPROVEMENT BEYOND THE STARTING POPULATION

Each of the six benchmark searches starts with 100 initial molecules. The single-parent antibiotic tasks are described separately in Appendix L. Table 11 compares the MolDeriveX initial population with its final exported set under the same scoring functions. The initial population and final output are evaluated separately.

Five of the six initial populations contain no joint hit; the dual-kinase population contains two. On the four-objective task, the initial mean scores are 0.5733 (QED), 0.7557 (SA), 0.4099 (GSK3B), and 0.2864 (JNK3). Thus, relatively favorable SA coexists with weak kinase predictions, leaving

substantial room for joint improvement. Optimization increases HV to 0.7103 and produces 381 joint hits.

At a matched size of 100 output records, mean HV over 100 sampled subsets is 0.5295, compared with 0.3597 for the initial 100 molecules. Recorded ancestor comparisons provide a complementary molecule-level assessment.

Table 11: Initial and final MolDeriveX sets, evaluated with the common scoring pipeline.

Task	$\|I\|$	Initial HV	Final HV	Initial hits	Final hits
QED–logP	100	0.2498	0.9470	0	2747
QED–SA	100	0.5925	0.9483	0	728
GSK3B–JNK3	100	0.6948	0.8942	2	1217
QED–SA–GSK3B	100	0.5688	0.8779	0	745
QED–SA–JNK3	100	0.5614	0.7694	0	315
QED–SA–GSK3B–JNK3	100	0.3597	0.7103	0	381

Table 12: Changes in the joint Top-100 relative to recorded parents and initial ancestors.

Task	Reference	Mean gain	Improved (%)	Dominates (%)
QED–logP	Direct parent	+0.0082	97.0	10.0
QED–logP	Initial ancestor	+0.4641	100.0	100.0
QED–SA	Direct parent	+0.0129	44.0	31.0
QED–SA	Initial ancestor	+0.2857	100.0	59.0
GSK3B–JNK3	Direct parent	-0.0131	27.0	13.0
GSK3B–JNK3	Initial ancestor	+0.2509	100.0	45.0
QED–SA–GSK3B	Direct parent	+0.0243	49.0	20.0
QED–SA–GSK3B	Initial ancestor	+0.2605	100.0	30.0
QED–SA–JNK3	Direct parent	-0.0285	32.0	6.0
QED–SA–JNK3	Initial ancestor	+0.2018	100.0	32.0
QED–SA–GSK3B–JNK3	Direct parent	+0.0388	72.0	16.0
QED–SA–GSK3B–JNK3	Initial ancestor	+0.2837	100.0	55.0

For the four-objective Top-100, all products improve their mean score over their initial ancestors and 55% dominate them. The mean gain of 0.2837 is distributed across GSK3B (+0.4091), JNK3 (+0.4153), QED (+0.2413), and SA (+0.0689). This pattern is consistent with improving the initially weaker activities while retaining favorable SA. The proposal-level analysis below complements these final elite-candidate results.

No final MolDeriveX molecule in the four target activity tasks is an exact initial-reference identity. Excluding initial identities therefore leaves their HV and hit counts unchanged. This supports improvement beyond simply retaining the starting molecules.

I.1 Exploratory proposals versus final elite candidates

For a product x with recorded immediate parent $p(x)$, mean gain is $d^{-1}\sum_j[f_j(x) - f_j(p(x))]$. Dominance requires no objective loss and at least one gain, using tolerance 10^{-12}. Recorded immediate parents and initial ancestors are distinct references. On the four-objective task, 11.64% of proposals improve their direct parent's mean objective score (Table 13). Across all 5,000 recorded proposal–parent pairs per task, mean gain is negative. Population selection subsequently determines which proposals are retained. Table 13, together with the ancestor gains in Table 12, distinguishes exploratory proposal quality from the selected final outputs. Even the Top-100 have negative mean direct-parent gains on the dual-kinase and QED–SA–JNK3 tasks. The contrast between proposal-level changes and final ancestor gains highlights the distinct roles of exploration and population selection.

Table 13: MolDeriveX generated records relative to their recorded immediate parents (5,000 pairs per task). Mean gain averages normalized objective changes. Improvement denotes positive mean gain; dominance requires no objective loss and at least one strict gain (tolerance 10^{-12}).

Task	Mean gain	Improved (%)	Dominates (%)
QED–logP	-0.000297	43.76	29.48
QED–SA	-0.123658	14.92	7.94
GSK3B–JNK3	-0.151180	9.86	4.08
QED–SA–GSK3B	-0.090745	13.74	2.58
QED–SA–JNK3	-0.112126	5.30	1.20
QED–SA–GSK3B–JNK3	-0.104678	11.64	1.18

J STRUCTURAL PRESERVATION AND EXPLORATION

For each task, let I be its 100-molecule MolDeriveX initial population. We define proximity by $s_I(x) = \max_{z \in I} \mathrm{Tan}(\phi(x), \phi(z))$ using radius-2 Morgan fingerprints (Rogers & Hahn, 2010). The comparisons restrict each method's candidates to $s_I > 0.4$ or $s_I > 0.6$, then compute HV and joint-hit counts. Every method is assessed against the same MolDeriveX initial-reference set. Subset HV uses all eligible molecules; table entries give HV / unique joint hits.

The neighborhood comparisons reveal task-dependent differences (Tables 14 and 15). At $s_I > 0.6$, MolDeriveX has 477 dual-target activity hits, versus at most 77 for a baseline. On QED–SA–GSK3B, it has three hits versus ten for MARS; on the four-objective task, it has none versus two for RetMol. RetMol therefore finds more joint hits within this four-objective neighborhood. Across the complete candidate sets, the corresponding counts are 381 for MolDeriveX and 9 for RetMol (Table 6). These results distinguish success near the starting set from success over the full output: MolDeriveX's four-objective hits occur outside the $s_I > 0.6$ neighborhood.

These neighborhoods describe whole-molecule similarity to the common initial reference.

Table 14: Target activity task outcomes within initial-reference similarity $s_I > 0.4$.

Method	GSK3β + JNK3	QED + SA + GSK3β	QED + SA + JNK3	QED + SA + GSK3β + JNK3
MARS	0.7866 / 83	0.8136 / 19	0.7939 / 12	0.5688 / 3
InversionGNN	0.7050 / 184	0.5662 / 0	0.4293 / 0	0.2085 / 0
MolSearch	0.0000 / 0	0.0000 / 0	0.0000 / 0	0.0000 / 0
GA+DL	0.7315 / 141	0.5893 / 0	0.5374 / 0	0.3244 / 0
MLPS	0.3212 / 0	0.3739 / 0	0.3694 / 0	0.0902 / 0
Hgraph	0.7788 / 57	0.7886 / 22	0.6620 / 26	0.5425 / 2
HN-GFN	0.0014 / 0	0.2437 / 0	0.0000 / 0	0.3312 / 1
RetMol	0.6836 / 42	0.8065 / 43	0.6437 / 9	0.5673 / 5
MolDeriveX	0.8609 / 769	0.8547 / 244	0.7662 / 186	0.6437 / 5

Table 15: Target activity task outcomes within initial-reference similarity $s_I > 0.6$.

Method	GSK3β + JNK3	QED + SA + GSK3β	QED + SA + JNK3	QED + SA + GSK3β + JNK3
MARS	0.7734 / 60	0.7787 / 10	0.7776 / 6	0.5372 / 1
InversionGNN	0.6979 / 73	0.5493 / 0	0.4285 / 0	0.1995 / 0
MolSearch	0.0000 / 0	0.0000 / 0	0.0000 / 0	0.0000 / 0
GA+DL	0.7252 / 77	0.5650 / 0	0.4660 / 0	0.2982 / 0
MLPS	0.0000 / 0	0.0000 / 0	0.2671 / 0	0.0261 / 0
Hgraph	0.5172 / 1	0.2338 / 0	0.4010 / 1	0.2070 / 0
HN-GFN	0.0000 / 0	0.0000 / 0	0.0000 / 0	0.0000 / 0
RetMol	0.6051 / 8	0.4730 / 3	0.5106 / 1	0.3623 / 2
MolDeriveX	0.8345 / 477	0.6836 / 3	0.6648 / 15	0.4178 / 0

K SENSITIVITY TO OUTPUT SUBSAMPLING

For each task t, we draw 100 uniform subsets of raw output records, without replacement within a subset, before validity filtering and deduplication. A common subset size is used across methods: $n_t = \min\{5000, \min_m N_{m,t}\}$, giving 103 for QED–logP, 295 for QED–SA, and 100 for the other tasks. This diagnostic compares coverage at a common sampled-output size.

Tables 16 and 17 report mean HV and empirical 2.5–97.5% subset ranges. On the four-objective task, MolDeriveX reaches mean HV 0.5295, above every baseline and its 100-molecule initial HV of 0.3597. The QED–SA ranking can instead favor Hgraph, so endpoint advantages are task dependent. The ranges summarize variation among the 100 output subsets.

Table 16: Hypervolume under matched output subsampling (two-objective tasks).

Method	QED + logP	QED + SA	GSK3β + JNK3
MARS	0.8393 [0.7342, 0.9099]	0.9223 [0.8967, 0.9408]	0.6642 [0.6282, 0.7031]
InversionGNN	0.7070 [0.6484, 0.7307]	0.6553 [0.6047, 0.7037]	0.5839 [0.5185, 0.6368]
MolSearch	0.6611 [0.6611, 0.6611]	0.9458 [0.9458, 0.9458]	0.6948 [0.6948, 0.6948]
GA+DL	0.8096 [0.7439, 0.8939]	0.9214 [0.9011, 0.9447]	0.5242 [0.4254, 0.6206]
MLPS	0.9388 [0.9387, 0.9388]	0.9375 [0.9373, 0.9375]	0.3545 [0.3292, 0.3831]
Hgraph	0.9468 [0.9455, 0.9473]	0.9479 [0.9469, 0.9481]	0.6703 [0.5176, 0.7490]
HN-GFN	0.8518 [0.7468, 0.9357]	0.9033 [0.8650, 0.9279]	0.2580 [0.2129, 0.3137]
RetMol	0.9354 [0.9113, 0.9448]	0.9475 [0.9451, 0.9481]	0.6225 [0.5437, 0.6971]
MolDeriveX	0.9441 [0.9353, 0.9467]	0.9439 [0.9355, 0.9482]	0.7660 [0.7076, 0.8220]

Table 17: Hypervolume under matched output subsampling (three- and four-objective tasks).

Method	QED + SA + GSK3β	QED + SA + JNK3	QED + SA + GSK3β + JNK3
MARS	0.5870 [0.4599, 0.7016]	0.5212 [0.3834, 0.6829]	0.2679 [0.1755, 0.4107]
InversionGNN	0.3061 [0.2290, 0.4328]	0.2234 [0.1366, 0.3039]	0.0760 [0.0454, 0.1251]
MolSearch	0.5688 [0.5688, 0.5688]	0.5614 [0.5614, 0.5614]	0.3597 [0.3597, 0.3597]
GA+DL	0.2640 [0.1369, 0.4262]	0.2840 [0.1079, 0.4382]	0.1331 [0.0715, 0.2162]
MLPS	0.4786 [0.4561, 0.5243]	0.3526 [0.2730, 0.4683]	0.0462 [0.0128, 0.0841]
Hgraph	0.6707 [0.5025, 0.7711]	0.5734 [0.4217, 0.6479]	0.3519 [0.1654, 0.5153]
HN-GFN	0.1946 [0.1221, 0.2826]	0.1000 [0.0415, 0.1778]	0.0445 [0.0138, 0.2140]
RetMol	0.7018 [0.5827, 0.7859]	0.5178 [0.3834, 0.6780]	0.3243 [0.2160, 0.4975]
MolDeriveX	0.7771 [0.7288, 0.8209]	0.6343 [0.5789, 0.6880]	0.5295 [0.4730, 0.6025]

L ANTIBIOTIC DERIVATIZATION CASE STUDY

L.1 OBJECTIVES AND SCOPE

The case study comprises three settings, each initialized from one parent: linezolid, clindamycin, or erythromycin. Descendants expand the population during optimization. The same pretrained evaluators and editing resources are used across these settings. The objectives are a pretrained Chemprop message-passing activity score and structural similarity to the parent. We use the pretrained activity model supplied by Liu et al. (2023), developed for predicting growth inhibition of *Acinetobacter baumannii*. Chemprop learns graph-based representations for molecular property prediction (Yang et al., 2019). The activity score lies in $[0, 1]$, with larger values indicating higher predicted antibacterial activity.

For candidate x and parent p, let $b(x)$ and $b(p)$ be binary Morgan fingerprints with radius 2 and 2,048 bits. The similarity objective is

$$S_{\mathrm{sim}}(x, p) = \frac{|b(x) \cap b(p)|}{|b(x)| + |b(p)| - |b(x) \cap b(p)|}. \tag{20}$$

The two maximization objectives are $\mathbf{f}_p(x) = (S_{\mathrm{bio}}(x), S_{\mathrm{sim}}(x, p))$. Similarity is optimized as a soft objective alongside predicted activity. The case figures compare MolDeriveX with GA+D, MolSearch, and F-CReM. GA+D denotes the neural-guided genetic comparator (Nigam et al., 2020); MolSearch uses molecular transformations (Sun et al., 2022). F-CReM denotes the case-specific CReM comparator; CReM supplies context-dependent fragment mutations (Polishchuk, 2020).

L.2 AUXILIARY TOXICITY AND DRUG-LIKENESS ASSESSMENT

We trained a binary molecular property classifier with Chemprop (Yang et al., 2019) on the ClinTox dataset (Wu et al., 2018), using the ClinTox data distributed with Chemprop. The classifier output is used as the toxicity score for the initial parent and the displayed derivatives in Figure 4. Lower scores indicate lower model-predicted toxicity. This classifier is separate from the pretrained antibacterial activity model of Liu et al. (2023).

QED is calculated using RDKit (Bickerton et al., 2012) as a descriptor-based measure of drug-likeness. Toxicity and QED are auxiliary readouts of the displayed sequence; the two optimization objectives are antibacterial activity and parent-relative similarity.

L.3 OBJECTIVE-SPACE AND DERIVATIVE-REGION METRICS

For a scored candidate set S, two-dimensional HV is the area dominated relative to $(0, 0)$:

$$\mathrm{HV}(S) = \lambda_2 \left(\bigcup_{x \in S} [0, S_{\mathrm{bio}}(x)] \times [0, S_{\mathrm{sim}}(x, p)] \right). \tag{21}$$

Both axes retain their common $[0, 1]$ score scales. Larger HV reflects better joint coverage of predicted activity and parent-relative similarity. The ideal-point distance associated with the singleton reference $P^* = \{(1, 1)\}$ is

$$d_{\mathrm{ideal}}(S) = \min_{x \in S} \|\mathbf{f}_p(x) - (1, 1)\|_2. \tag{22}$$

The figure labels this singleton-reference distance as IGD. Smaller values indicate closer proximity to the ideal point $(1, 1)$.

Diversity of derivative regions. We measure the diversity of the structural modifications separately for each parent. We first exclude the fixed parent scaffold from each candidate, extract the remaining derivative regions, and remove duplicate region identities. Let $\mathcal{R}_p(S) = \{R_1, \ldots, R_n\}$ denote the resulting distinct derivative-region graphs, and let $\phi(R_i)$ be the extended-connectivity fingerprint (ECFP) of region R_i. Diversity is the mean Tanimoto distance over all unordered pairs of these fingerprints:

$$D_{\text{deriv}}(S \mid p) = \frac{2}{n(n-1)} \sum_{i<j} \left[1 - \text{Tan}\big(\phi(R_i), \phi(R_j)\big)\right], \qquad n \geq 2. \tag{23}$$

Here n is the number of unique derivative regions. The score lies in $[0, 1]$; larger values indicate more varied modifications. Computing fingerprints after excluding the shared scaffold focuses the comparison on the derived portions instead of the common parent structure. It differs from whole-molecule internal diversity $D(U)$ in Appendix E, whose fingerprints include the complete candidate. The two metrics characterize different levels of structural variation.

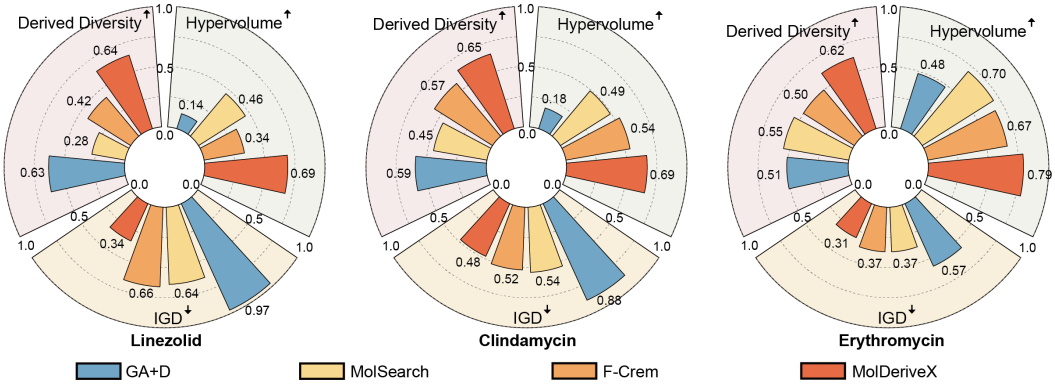


Figure 5: Quantitative summary for three single-parent antibiotic tasks, in the order linezolid, clindamycin, and erythromycin. Larger HV and derivative-region diversity and smaller ideal-point distance (IGD) indicate better performance. Derived diversity averages pairwise ECFP Tanimoto distances over unique derivative regions after excluding the fixed parent scaffold.

Starting from one parent per task, the reported HV values for MolDeriveX are 0.69, 0.69, and 0.79; the strongest comparator values are 0.46, 0.54, and 0.70, respectively. Derivative-region diversity is 0.64, 0.65, and 0.62, compared with the strongest comparator values of 0.63, 0.59, and 0.55. The diversity gains are 0.01, 0.06, and 0.07, respectively. The figure also reports smaller distance values for MolDeriveX (0.34, 0.48, and 0.31).

L.4 Property changes along the structural sequence

Figure 4 follows a linezolid derivatization sequence. Toxicity and QED are evaluated as described in Appendix L.2. The score changes show that larger predicted activity can accompany a limited shared structural region and lower whole-molecule similarity. The toxicity curve falls most strongly at the first step, then remains near its lower value with a small final increase; QED is modestly lower than the parent's. The case evaluates computational derivatization through predicted *A. baumannii* activity and parent similarity.

M CANDIDATE DIVERSITY, UNIQUENESS, AND NOVELTY

Table 18 complements objective quality with three whole-output metrics. Internal diversity $D(U)$ is defined in Appendix E. Uniqueness is $|U|/V$, where V counts valid records before deduplication. Molecular novelty is the fraction of U absent by canonical identity from the common ZINC250k reference of 249,428 unique molecules. Novelty is defined relative to this common reference database. All metrics use the full candidate set.

MolDeriveX has 100% uniqueness and 99.96–100% novelty across the six tasks, with diversity ranging from 0.7702 to 0.8583. On the four-objective task, diversity is 0.7897 compared with Hgraph's 0.8738. These results distinguish generating distinct, reference-novel candidates from maximizing structural dispersion; the method's strongest advantage remains joint objective quality and successful structural coverage.

Table 18: Overall candidate diversity, uniqueness, and reference-relative novelty across all six tasks.

Method	Metric	QED + logP	QED + SA	GSK3β + JNK3	QED + SA + GSK3β	QED + SA + JNK3	QED + SA + GSK3β + JNK3
MARS	$D(U)$	0.9096	0.8959	0.8721	0.8939	0.8926	0.8923
	Unique (%)	73.00	79.68	71.28	71.06	72.48	73.32
	Novel (%)	99.97	100.00	100.00	99.94	99.92	99.95
InversionGNN	$D(U)$	0.7970	0.8477	0.8329	0.7767	0.7844	0.7624
	Unique (%)	100.00	100.00	100.00	100.00	100.00	100.00
	Novel (%)	100.00	100.00	100.00	100.00	100.00	100.00
MolSearch	$D(U)$	0.9082	0.8401	0.8674	0.8759	0.8697	0.8741
	Unique (%)	100.00	100.00	100.00	100.00	100.00	100.00
	Novel (%)	100.00	99.66	100.00	100.00	100.00	100.00
GA+DL	$D(U)$	0.9038	0.9079	0.8819	0.8685	0.8915	0.8890
	Unique (%)	100.00	100.00	100.00	100.00	100.00	100.00
	Novel (%)	100.00	100.00	100.00	100.00	100.00	100.00
MLPS	$D(U)$	0.8908	0.8892	0.8899	0.8757	0.9058	0.9058
	Unique (%)	37.50	30.17	78.14	76.28	91.78	93.07
	Novel (%)	100.00	100.00	100.00	100.00	100.00	100.00
Hgraph	$D(U)$	0.8384	0.8412	0.8701	0.8631	0.8700	0.8738
	Unique (%)	82.35	61.69	95.30	84.78	85.32	92.42
	Novel (%)	99.71	99.51	99.96	99.73	99.78	99.91
HN-GFN	$D(U)$	0.8105	0.9098	0.7989	0.9038	0.9135	0.9064
	Unique (%)	84.54	93.74	99.98	95.92	93.04	94.08
	Novel (%)	100.00	99.98	100.00	100.00	100.00	99.98
RetMol	$D(U)$	0.8355	0.8153	0.8427	0.8472	0.8243	0.8522
	Unique (%)	77.64	76.56	72.30	72.36	81.74	73.44
	Novel (%)	100.00	99.90	99.94	99.81	99.83	99.89
MolDeriveX	$D(U)$	0.8459	0.8583	0.8481	0.8367	0.7702	0.7897
	Unique (%)	100.00	100.00	100.00	100.00	100.00	100.00
	Novel (%)	100.00	99.96	99.98	99.96	100.00	99.96

M.1 WITHIN-METHOD CHEMICAL-SPACE VISUALIZATION

Figure 6 displays the four-objective candidate sets using t-SNE (Van der Maaten & Hinton, 2008). Each method is embedded independently from binary Tanimoto distances between 2,048-bit ECFP6-style Morgan fingerprints (radius 3, chirality disabled). The first 5,000 output rows are considered in source order, or all rows when fewer are available, without score-based selection. Invalid molecules are excluded; repeated valid molecules remain in the fit. The visualization uses perplexity 30, random initialization, an automatic learning rate, and 1,000 iterations. This radius-3 representation is specific to the figure; quantitative diversity in Table 18 uses radius-2 fingerprints.

Colors encode the mean normalized objective score on a common $[0, 1]$ scale. MolDeriveX places high-scoring candidates in multiple visible neighborhoods. Each panel has independently fitted coordinates; quantitative comparisons of structural diversity and scaffold coverage use the tables.

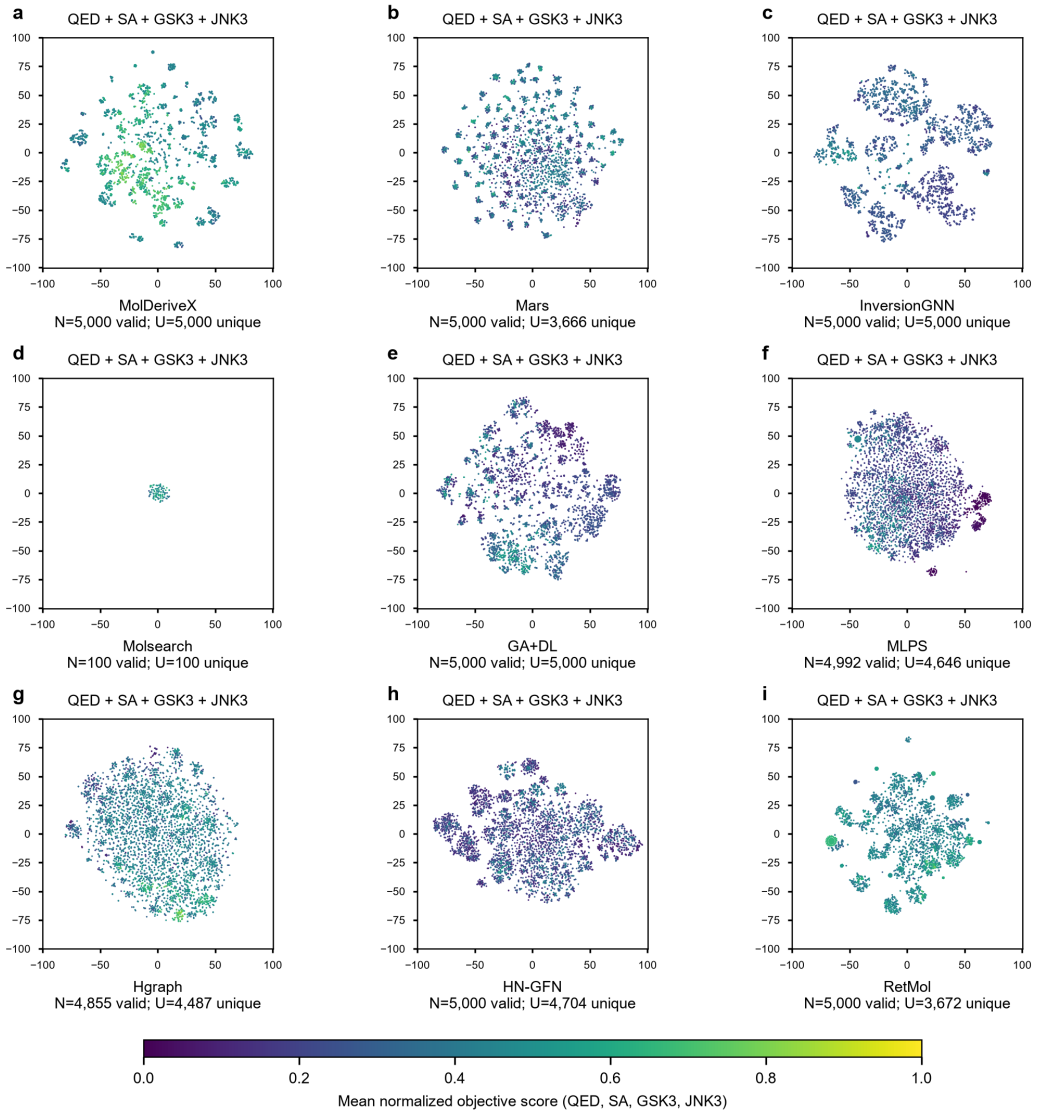


Figure 6: Within-method t-SNE visualization on the four-objective task. Panels use independently fitted ECFP6-style fingerprint embeddings. Colors denote mean normalized objective scores; N and U count included valid records and unique structures. Repeated valid molecules are retained.

N STRUCTURAL SCREENING OF SUCCESSFUL CANDIDATES

For the successful subset, novelty uses the same reference as Appendix M, with joint hits as its denominator. Combined screening requires one connected component and no RDKit PAINS or Brenk alert RDKit contributors (b). Screening percentages use joint hits as the denominator; a dash indicates an empty hit set. Neither screen is used to remove candidates from the primary HV or hit analysis.

MolDeriveX's high hit-set novelty (at least 99.7% across tasks) coexists with screening-pass fractions of 46.3–83.7%. Table 19 retains screening results for every method, including baselines with higher pass fractions.

Table 19: Combined structural-screening pass percentages among joint hits.

Method	QED + logP	QED + SA	GSK3β + JNK3	QED + SA + GSK3β	QED + SA + JNK3	QED + SA + GSK3β + JNK3
MARS	57.3	71.3	61.1	59.3	71.4	100.0
InversionGNN	85.2	—	33.8	—	—	—
MolSearch	100.0	95.9	50.0	—	—	—
GA+DL	10.7	54.2	16.0	—	—	—
MLPS	41.5	59.3	—	11.1	—	—
Hgraph	83.4	78.5	45.1	95.8	88.8	95.5
HN-GFN	74.6	86.5	0.0	—	—	100.0
RetMol	75.1	89.3	95.9	55.0	93.8	100.0
MolDeriveX	75.5	83.7	46.3	71.8	52.1	82.7

O ACTIVITY TRADE-OFFS AND REMAINING BOTTLENECKS

Figure 7 directly displays nondominated objective vectors in all three two-objective tasks. The dual-kinase frontier illustrates the trade-off coverage behind the HV advantage, while the descriptor-only panels show different method rankings. All methods are shown on the same normalized objective axes.

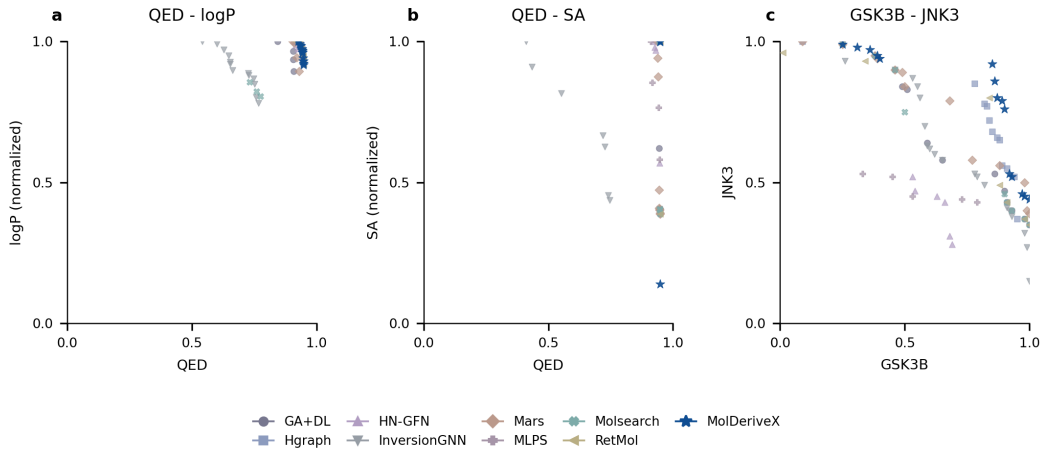


Figure 7: Two-objective nondominated score vectors from the full scored candidate sets. Duplicate vectors are collapsed; axes use the fixed normalized $[0, 1]$ scales.

Figure 8 separates activity gains from residual objective failures. On the four-objective task, MolDeriveX has lower marginal failure fractions for both kinase objectives than the baselines, while QED remains below its threshold for 72.98% of its candidates. QED therefore remains a major bottleneck despite improved activity coverage.

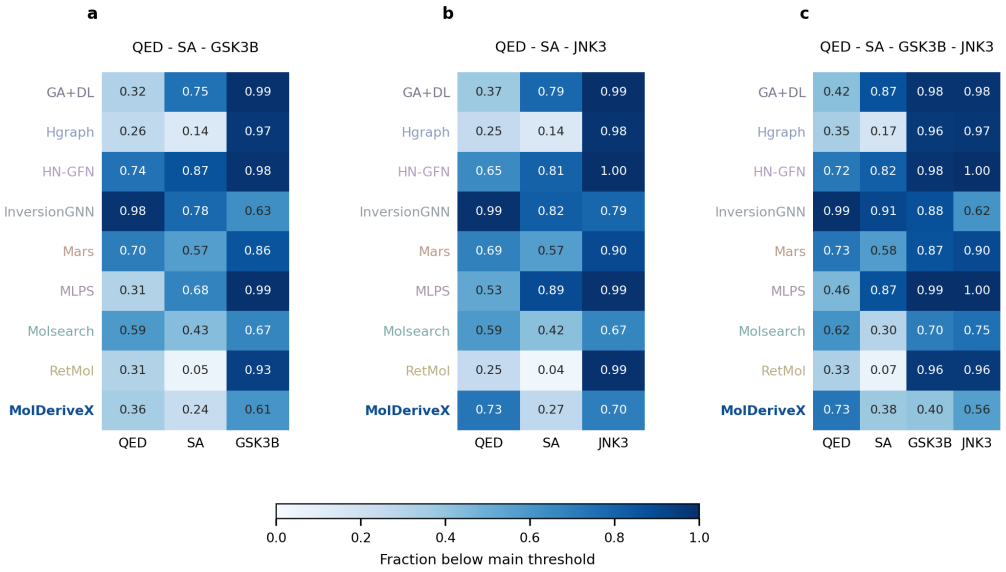


Figure 8: Objective-specific failure fractions in the three- and four-objective tasks. Each cell gives the fraction of scored unique molecules below that objective's main threshold. Darker cells indicate larger fractions. Fractions are computed separately for each objective; missing values are shown explicitly.

P Ablation of optimization quality and structural coverage

Table 3 reports HV on the five main-text tasks. Table 20 adds QED–logP and structural metrics for Full, local-only anchoring, removal of structural retention, and removal of lineage rescue across all six tasks. HV uses the exact zero-reference computation. To assess structural exploration alongside objective quality, we report overall internal diversity $D(U)$ and two complementary scaffold summaries on the valid unique candidate set U, including both successful and other scored molecules.

$D(U)$ averages dissimilarity over all unordered molecular pairs using the radius-2, 2,048-bit Morgan fingerprints defined in Appendix E. K_U counts distinct nonempty Bemis–Murcko scaffolds, while $K_{\text{eff}} = \exp(-\sum_s p_s \log p_s)$ measures their effective number, with p_s normalized over scaffold-bearing molecules. The former measures richness; the latter also accounts for concentration. The table includes $|U|$ because scaffold richness depends on the number of distinct molecules. Together, these summaries characterize the structural coverage of the evaluated candidate sets.

Table 20: Optimization quality and structural coverage of the completed ablation configurations.

Task	Configuration	$\lvert U\rvert$	HV ↑	$D(U)$ ↑	K_U ↑	K_{eff} ↑
QED–logP	Full	5000	**0.9456**	0.8550	3646	1978.1
	Local-only	5000	0.9030	0.8632	**3881**	**2937.7**
	w/o structural	5000	0.8968	**0.8679**	3849	2882.0
	w/o lineage	5000	0.9272	0.8618	3839	2826.6
QED–SA	Full	5000	**0.9483**	0.8583	2957	1385.6
	Local-only	5000	0.9418	**0.8708**	**3177**	**1611.0**
	w/o structural	5000	0.9461	0.8653	2943	1360.7
	w/o lineage	5000	0.9444	0.8702	2772	1329.4
GSK3B–JNK3	Full	5000	**0.8942**	**0.8481**	2972	1622.5
	Local-only	5000	0.8476	0.8423	**3228**	**1912.3**
	w/o structural	4722	0.8645	0.8412	2916	1586.8
	w/o lineage	4859	0.7981	0.7848	2761	1616.1
QED–SA–GSK3B	Full	5000	**0.8779**	0.8367	**2889**	1455.8
	Local-only	5000	0.7878	0.8376	2625	1379.8
	w/o structural	5000	0.8202	0.8461	2869	**1721.5**
	w/o lineage	5000	0.8505	**0.8476**	2542	1290.2
QED–SA–JNK3	Full	5000	**0.7694**	0.7702	1748	667.1
	Local-only	5000	0.7440	**0.7743**	1996	758.8
	w/o structural	5000	0.7348	0.7570	**2793**	**1519.2**
	w/o lineage	5000	0.7386	0.7518	1894	682.3
QED–SA–GSK3B–JNK3	Full	5000	**0.7103**	0.7897	**3697**	2374.4
	Local-only	5000	0.6213	0.7907	3290	1900.7
	w/o structural	5000	0.6662	0.7842	3671	**2405.0**
	w/o lineage	5000	0.6600	**0.7940**	3292	1858.7

Full has the highest HV in all six configuration comparisons. On the four-objective task it combines HV 0.7103 with the largest scaffold count (3,697), compared with 3,290 for local-only anchoring and 3,292 without lineage rescue. Its effective scaffold count is 2,374.4, versus 1,900.7 and 1,858.7, respectively. This supports improved objective quality together with broad scaffold coverage. Removing lineage rescue gives slightly higher fingerprint diversity (0.7940 versus 0.7897), and removing structural retention gives a slightly higher effective scaffold count (2,405.0). The configurations therefore differ in their balance of objective quality, fingerprint diversity, and scaffold coverage.

The dual-kinase task provides a complementary pattern: without lineage rescue, diversity falls from 0.8481 to 0.7848 and HV falls from 0.8942 to 0.7981. However, structural rankings vary across tasks. Together, the results support the complete configuration's balance of objective quality and structural coverage.