
Directly Optimizing for Synthesizability in Generative Molecular Design using Retrosynthesis Models

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Synthesizability in generative molecular design remains a pressing challenge.
2 Existing methods to assess synthesizability span heuristics-based methods, ret-
3 rosynthesis models, and synthesizability-constrained molecular generation. The
4 latter has become increasingly prevalent and proceeds by defining a set of permitted
5 actions a model can take when generating molecules, such that all generations are
6 anchored in "synthetically-feasible" chemical transformations. To date, retrosynthe-
7 sis models have been mostly used as a post-hoc filtering tool as their inference cost
8 remains prohibitive to use directly in an optimization loop. In this work, we show
9 that with a sufficiently sample-efficient generative model, it is straightforward to
10 directly optimize for synthesizability using retrosynthesis models in goal-directed
11 generation. Under a heavily-constrained computational budget, our model can
12 generate molecules satisfying a multi-parameter drug discovery optimization task
13 while being synthesizable, as deemed by the retrosynthesis model.

14 1 Introduction

15 Generative molecular design for drug discovery has recently seen a surge of experimental validation,
16 with many candidate molecules progressing into clinical trials¹. However, the synthesizability of
17 generated designs remains a pressing challenge. Regardless of how "good" generated molecules are,
18 they must be synthesized and experimentally validated to be of use, and work has shown that many
19 generative models propose molecules for which finding a viable synthetic route for, is *at the very*
20 *least* not straightforward^{2,3}. Existing works tackle synthesizability in generative molecular design
21 either by heuristics^{4,5}, learning synthetic complexity from reaction corpus⁶, retrosynthesis models
22 which predict synthetic routes⁷⁻¹⁸, or enforce a notion of synthesizability directly in the generative
23 process¹⁹⁻²⁹.

24 Recently, synthesizability-constrained generative models¹⁹⁻²⁹ have become increasingly prevalent.
25 A typical metric to *quantify* synthesizability is whether a retrosynthesis model can solve a route for
26 the generated molecules^{28,29}. It is common practice to apply retrosynthesis models during post-hoc
27 filtering due to their inference cost^{2,3}.

28 On the other hand, sample efficiency is also a pressing challenge, which concerns with how many
29 oracle calls (computational predictions of molecular properties) are required to optimize an objective
30 function. When these oracle calls are computationally expensive, such as in binding affinity pre-
31 dictions, there is a practical limit to an *acceptable* oracle budget for real-world model deployment.
32 The Practical Molecular Optimization (PMO) benchmark³⁰ highlighted the importance of sample
33 efficiency and since then, more recent works have explicitly considered an oracle budget³¹⁻⁴⁰.

34 Recently, Saturn⁴⁰, which is a language-based molecular generative model leveraging the Mamba⁴¹
35 architecture, displayed state-of-the-art sample efficiency compared to 22 models. In this work, we

36 build on Saturn and show that with a *sufficiently sample-efficient* model, one can treat retrosyn-
37 thesis models as an oracle⁴² and directly optimize for generating molecules where synthesis routes
38 can be solved for (Fig. 1). We compare to the recent Reaction-GFlowNet (RGFN)²⁹ model and
39 show that Saturn can optimize their proposed multi-parameter optimization (MPO) task to generate
40 molecules with good docking scores (to predict binding affinity) and is synthesizable (as deemed by
41 a retrosynthesis model) with 1/400th the oracle budget (1,000 calls instead of 400,000).

42 **We strictly emphasize that we neither claim to solve synthesizability nor claim our model**
43 **guarantees synthesizability.** Rather, the take-home message of this work is that if one wants
44 to optimize certain properties, then they should be included in the MPO objective function. If the
45 downstream metric is whether a retrosynthesis tool can solve a route for the generated molecules^{2,28,29},
46 then the tool itself should be part of the objective function (given that the model is not synthesizability-
47 constrained).

48 2 Related Work

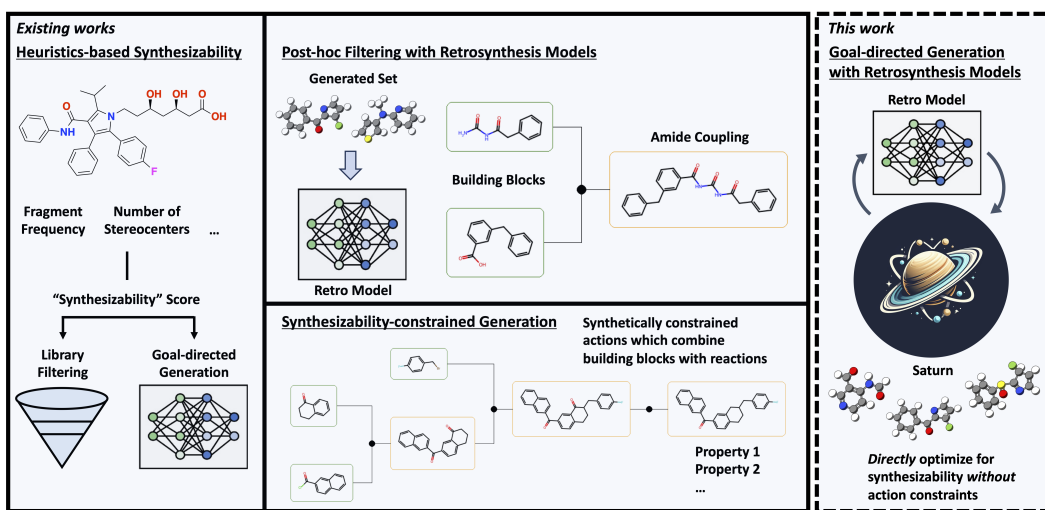


Figure 1: Overview of algorithmic methods to handle synthesizability in generative molecular design.

49 **Synthesizability Metrics.** Quantifying and defining synthesizability is non-trivial and early metrics
50 assess *molecular complexity* rather than synthesizability explicitly. Exemplary works include the
51 Synthetic Accessibility (SA) score⁴ and SYnthetic Bayesian Accessibility (SYBA)⁵ which are
52 based on the frequency of chemical groups in databases. The Synthetic Complexity (SC) score⁶ is
53 trained on Reaxys data to measure molecular complexity and implicitly considers the number of
54 synthetic steps required to make a target molecule. There is a correlation between these scores and
55 whether retrosynthesis tools can solve a route⁴³. The recent Focused Synthesizability (FS) score⁴⁴
56 incorporated domain-expert *preferences*⁴⁵ to assess synthesizability.

57 **Retrosynthesis Models.** Given a target molecule, retrosynthesis models propose viable synthetic
58 routes by combining commercial building blocks (starting reagents) with reaction templates (coded
59 patterns that map chemical reaction compatibility) or template-free approaches (learned patterns from
60 data). Exemplary examples include the first work applying Monte Carlo tree search (MCTS) for
61 retrosynthesis¹⁰, SYNTHIA^{46,47}, AiZynthFinder¹¹⁻¹³, ASKCOS¹⁵, Eli Lilly’s LillyMol retrosynthe-
62 sis model⁴⁸, Molecule.one’s M1 platform⁴⁹, and IBM RXN^{16,50,51}. We further highlight surrogate
63 models including Retrosynthesis Accessibility (RA) score¹⁴ and RetroGNN¹⁷ trained on the output of
64 retrosynthesis models for faster inference. Note that these models output a score rather than synthetic
65 routes.

66 **Synthesizability-constrained Molecular Generation.** More recently, molecular generative models
67 have been designed with a notion of synthesizability, for example by enforcing transformations
68 from a set of permitted reaction templates. Expansion methods include SYNOPSIS¹⁹, Design
69 of Genuine Structures (DOGS)²⁰, and RENATE⁵². Other models include MOLECULE CHEF²¹,

70 Synthesis Directed Acyclic Graph (DAG)²², ChemBO²³, SynNet²⁶, and SyntheMol²⁷. Models
71 that also use reinforcement learning (RL) include Policy Gradient for Forward Synthesis (PGFS)²⁴,
72 Reaction-driven Objective Reinforcement (REACTOR)²⁵, and LibINVENT⁵³. Recent works have
73 equipped GFlowNets⁵⁴ with reaction templates, including SynFlowNet²⁸ and RGFN²⁹. Finally, very
74 recent work proposes a new paradigm of "projecting" unsynthesizable molecules into similar, but
75 synthesizable analogs⁵⁵.

76 **Goal-directed Generation with Synthesizability Metrics.** An alternative to synthesizability-
77 constrained molecular generation is to task molecular generative models to also optimize for synthe-
78 sizability metrics⁴⁴, with common ones being SA score^{2,3}. Although SA score assesses molecular
79 complexity, it is correlated with whether AiZynthFinder can solve a route⁴³. Generally, more con-
80 fidence is placed on the output of retrosynthesis models in assessing synthesizability and this is
81 reflected in works that assess model performance on whether generated molecules have a solved
82 route^{28,29,55}. In this work, we propose to directly incorporate retrosynthesis tools as an oracle in the
83 MPO objective function and show that generated molecules satisfy MPO objectives.

84 We end this section by reinforcing that quantifying synthesizability is non-trivial and neither reaction
85 templates nor retrosynthesis tools *guarantee* synthesizability. Notably, reaction templates depend on
86 the granularity of their definition, for instance with the inclusion or omission of incompatibilities
87 which affects the false positive rate of matching reagents^{56,57}. A concrete example of this is the
88 original paper reporting Enamine REAL which is a "make-on-demand" commercial database with
89 a stated ~80% synthesis success rate⁵⁸. Recently, SyntheMol²⁷ which enforces reaction templates
90 during molecular generation, ordered 70 compounds from Enamine REAL with 58 successful
91 syntheses (~83%). We wish to emphasize that the point of drawing attention to this is strictly to
92 support our statement that neither reaction templates, make-on-demand libraries (often generated by
93 reaction templates), nor retrosynthesis tools *guarantee* synthesizability. "Make-on-demand libraries"
94 are a remarkable resource.

95 3 Methods

96 In this section, we describe in detail the experimental design and highlight caveats in the results.
97 Firstly, we use Saturn⁴⁰ as the generative model which uses RL for goal-directed generation and has
98 high sample efficiency. For details of the model, we refer to the original work⁴⁰. In Saturn, we newly
99 implement AiZynthFinder¹¹⁻¹³ and QuickVina2-GPU-2.1⁵⁹⁻⁶¹ (for docking) as oracles to match the
100 case study in RGFN²⁹.

101 The RGFN work assesses the synthesizability of generated molecules using the quantitative estimate
102 of drug-likeness (QED)⁶², SA score⁴, and whether AiZynthFinder¹¹⁻¹³ can solve a route. The authors
103 state that the latter better estimates synthesizability. This statement is supported by recent work
104 highlighting how AiZynthFinder predictions can augment medicinal chemists' decision-making,
105 which led to real-world impact in commercial drug discovery projects⁶³. The RGFN work features
106 three case studies where the objective function is either to optimize a proxy model for docking scores,
107 optimize a proxy model for biological activity classification, or optimize QuickVina2-GPU-2.1⁵⁹⁻⁶¹
108 docking scores directly. We choose to compare our model on the latter task because proxy models,
109 while offering faster inference, suffer from domain out-of-applicability if generated molecules deviate
110 too far from the training data. This was also stated by the authors and was the motivation for designing
111 the docking case study²⁹.

112 **Experimental Caveats.** As the code for RGFN²⁹ is not released, we implement their oracle function
113 ourselves. In Appendix C, we describe the steps we took to reproduce their case study faithfully. Here,
114 we instead highlight caveats that make the comparison not exactly apples to apples for transparency.

- 115 1. **Pre-training:** Saturn is pre-trained with either ChEMBL 33⁶⁴ or ZINC⁶⁵. These datasets,
116 containing bio-active molecules, inherently bias the learned distribution to already known
117 synthesizable entities². On the other hand, RGFN defines a state space based on reaction
118 templates and building blocks. We note, however, that these are common pre-training
119 datasets that many generative models in literature are pre-trained with.
- 120 2. **Quantifying Synthesizability:** RGFN handles synthesizability by combining building
121 blocks with reaction templates. By contrast, Saturn is handling synthesizability by op-
122 timizing AiZynthFinder. It could be the case that RGFN's reaction templates represent

123 "true" synthesizability better in some cases. We note, however, that RGFN also evaluates
124 synthesizability using AiZynthFinder, and in principle, any building blocks and templates
125 used in RGFN could be added to a retrosynthesis model.

126 3. **Docking Results Filtering.** RGFN filters the best generated molecules by whether they pass
127 PoseBusters⁶⁶ checks (plausible physicality). We do not consider this as this is essentially an
128 artefact of the oracle. Provided a more accurate oracle, failure naturally decreases. We note
129 that QuickVina2-GPU-2.1 outputs almost all pass the PoseBusters checks (see Appendix
130 H in the RGFN²⁹ work). Subsequently, RGFN filters molecules to be dissimilar (< 0.4
131 Tanimoto similarity) to known aggregators based on the Aggregation Advisor dataset⁶⁷
132 as the *final set*. We also do not consider this. RGFN reports statistics *before* this final
133 aggregator filtering and *these* are the results we compare to (Table 1 in the RGFN²⁹ work).

134 4. **Objective Function.** The RGFN work (which also reports results for GraphGA⁶⁸, Synthe-
135 Mol²⁷, and FGFN⁶⁹), defines the objective function to only optimize for docking score, but
136 assesses generated molecules also by their QED and SA scores. It is unclear the performance
137 of these models if the objective function were modified to also enforce these properties.

138 Still, despite these caveats, the message we convey is that if one wants to optimize for downstream
139 metrics, then they should be included in the MPO objective function. This is often impractical
140 because certain oracles are computationally expensive and generative models are not efficient enough
141 to directly optimize them. Provided a model *is* sufficiently sample-efficient, generative models can be
142 tasked to optimize *anything* (this does *not* mean that it will *always* be able to optimize the objective
143 under the budget).

144 **Experimental Setup.** Following the RGFN²⁹ work, the case study is to generate synthesizable
145 molecules with good docking scores (using QuickVina2-GPU-2.1⁵⁹⁻⁶¹) to ATP-dependent Clp
146 protease proteolytic subunit (ClpP). The objective function is:

$$R_{RGFN}(x) = \text{Docking Score}(x) \quad (1)$$

147 where x is a generated molecule. In Saturn, we apply reward shaping so that $R_{RGFN}(x) \in [0, 1]$. As
148 the purpose of this short paper is to convey that retrosynthesis models can be directly optimized as an
149 oracle, we further define two objective functions:

$$R_{\text{All MPO}}(x) = (\text{Docking Score}(x) \times \text{QED}(x) \times \text{SA Score}(x) \times \text{AiZynthFinder}(x))^{\frac{1}{4}} \in [0, 1] \quad (2)$$

$$R_{\text{Double MPO}}(x) = (\text{Docking Score}(x) \times \text{AiZynthFinder}(x))^{\frac{1}{2}} \in [0, 1] \quad (3)$$

150 See Appendix H for reward shaping details to normalize Eq. 2 and $3 \in [0, 1]$ and the exponential
151 term which is from the product aggregator that outputs the final reward. The rationale for $R_{\text{All MPO}}$
152 (Eq. 2) is because RGFN evaluates generated molecules also by their QED, SA score, and whether
153 AiZynthFinder can solve a route. Since these are the downstream metrics, we include them in the
154 objective function. The rationale for $R_{\text{Double MPO}}$ is to illustrate a contrast in optimization difficulty
155 as $R_{\text{All MPO}}$ is inherently more challenging. Still, we show in the Results section that both objective
156 functions can be optimized.

157 All Saturn experiments are run across 10 seeds (0-9 inclusive) with 1,000 oracle calls. We note this is
158 1/400th of the oracle budget of the RGFN work (400,000 calls). We compare with RGFN and also
159 GraphGA⁶⁸, SyntheMol²⁷, and Fragment-based GFlowNet (FGFN)⁶⁹. We do not run these models
160 ourselves and take the results from the RGFN work.

161 **Metrics.** Following the RGFN²⁹ work, a **Mode** is defined as a molecule with docking score < -10 .
162 **Discovered Modes** denotes the set of generated Modes that also possess Tanimoto similarity < 0.5
163 to every other mode. We note that Modes with > 0.5 Tanimoto similarity with other Modes are still
164 valuable, as given a pair of "similar" molecules, there can be a clear preference if for example, one of
165 the molecules contains an undesired substructure. For Saturn results, we additionally report **Yield**
166 which denotes the total number of unique molecules generated with docking score < -10 .

167 4 Results and Discussion

168 We devise three experiments: optimizing only docking score (following the RGFN²⁹ work), jointly
169 optimizing docking and AiZynthFinder, and lastly, showing that AiZynthFinder can *still* be directly
170 optimized as an oracle even if none of the molecules in the training data for the generative model can
171 be solved by AiZynthFinder. We construct a custom dataset for this last case study.

172 4.1 Experiment 1: Optimizing only docking score leads to unreasonable molecules

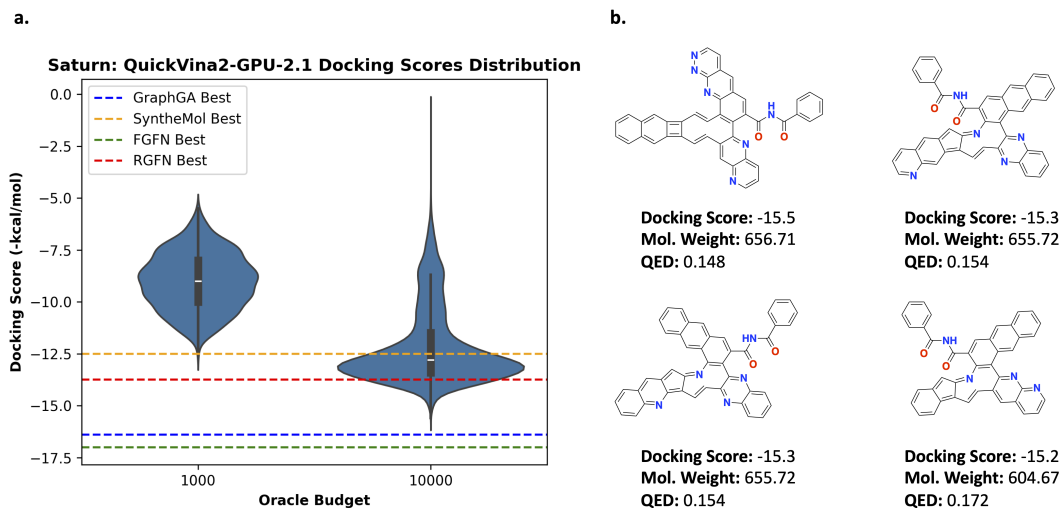


Figure 2: Experiment 1: Optimizing only docking scores. **a.** Distribution of docking scores at varying oracle budgets. The best docking score across comparison methods (taken from the RGFN²⁹ work) are annotated as dotted lines. **b.** Example lipophilic molecules generated by Saturn with the best docking scores.

173 We first present results for the R_{RGFN} (Eq. 1) objective function which only optimizes for docking
174 scores against ClpP. It is generally not advised to optimize this in isolation because docking oracles
175 can be highly exploitable, such that lipophilic (lots of carbon atoms and high logP) molecules
176 (promiscuous binders with solubility issues⁷⁰) receive good docking scores. We show that with
177 10,000 oracle calls, Saturn (trained on ChEMBL 33⁶⁴) generates molecules with approximately
178 the same best QuickVina2-GPU-2.1⁵⁹⁻⁶¹ docking scores compared to GraphGA⁶⁸, SyntheMol²⁷,
179 FGFN⁶⁹, and RGFN²⁹ which were run with 400,000 oracle calls (40x higher budget). We perform
180 one replicate here as we only want to convey that the objective function is highly exploitable. Fig. 2a
181 shows the distribution of docking scores at varying oracle budgets. We illustrate how the docking
182 oracle can be exploited in Fig. 2b which shows the best molecules generated by Saturn. Although
183 possessing good docking scores, they are lipophilic with high molecular weight and low QED.
184 Consequently, these are not meaningful molecules. Table 1 in the RGFN²⁹ work shows that the
185 best generated molecules across various models also have low QED: GraphGA (~0.32), FGFN
186 (~0.22), and RGFN (~0.23), suggesting that they are also exploiting the docking oracle. We note that
187 SyntheMol has slightly higher QED (~0.45).

188 4.2 Experiment 2: Directly optimizing for synthesizability using AiZynthFinder

189 In the previous section, we have shown that generative models can exploit docking oracles. Yet,
190 docking scores can be valuable as they can be *correlated* with better binding affinity⁷¹ and should be
191 optimized in combination with oracles that modulate physico-chemical properties. In this section,
192 we run Saturn with the $R_{All\ MPO}$ (jointly maximize QED, minimize SA score, minimize docking
193 score, and is AiZynthFinder solvable) and $R_{Double\ MPO}$ (jointly minimize docking score and is
194 AiZynthFinder solvable) objective functions.

195 **Quantitative Results.** Table 1 shows the Saturn results and also results taken from RGFN's²⁹
196 work. **As stated previously, since the comparison to RGFN is not apples to apples, we focus our**

Table 1: Synthesizability metrics for top-k Modes (**molecules with docking score < -10**). Results are taken from the RGFN²⁹ paper (it was not stated how many replicates the models were run for). Mol. weight, QED, and SA score results are for the top-500 Modes. AiZynth results are for the top-100 Modes. NR denotes "not reported". All Saturn experiments were run across 10 seeds (0-9 inclusive). The mean and standard deviation are reported. Both Yield and Modes are reported. The number after the configuration denotes the number of successful replicates out of 10 (Modes ≥ 1). For Saturn, none of the configurations found 100 Modes in 1,000 oracle calls so the metrics are reported for however many Modes were found.

^a Less than 400,000 oracle calls as SyntheMol²⁷ roll-outs took time. The RGFN authors decided to match the wall time instead.

Method	Modes (Yield)	Mol. weight (\downarrow)	QED (\uparrow)	SA score (\downarrow)	AiZynth (\uparrow)	Oracle calls (Wall time)
Previous work		top-500	top-500	top-500	top-100	
GraphGA ⁶⁸	NR	521.0 \pm 31.8	0.32 \pm 0.07	4.14 \pm 0.51	0.00	400,000 (NR)
SyntheMol ²⁷	NR	458.2 \pm 60.7	0.45 \pm 0.16	2.86 \pm 0.56	0.56	100,000 ^a (72h)
FGFN ⁶⁹	NR	548.6 \pm 42.9	0.22 \pm 0.03	2.94 \pm 0.54	0.25	400,000 (NR)
RGFN ²⁹	NR	526.2 \pm 37.6	0.23 \pm 0.04	2.83 \pm 0.22	0.65	400,000 (72h)
$R_{All\ MPO}$ (ours)		4 objectives (Docking, QED, SA, AiZynth)				
Saturn-ChEMBL (10)	4 \pm 1 (5 \pm 3)	367.7 \pm 15.7	0.70 \pm 0.13	2.11 \pm 0.19	0.91 \pm 0.11	1,000 (2.9h \pm 34m)
Saturn-GA-ChEMBL (10)	7 \pm 6 (10 \pm 9)	373.3 \pm 20.9	0.67 \pm 0.09	2.08 \pm 0.23	0.82 \pm 0.17	1,000 (2.1h \pm 24m)
Saturn-ZINC (9)	6 \pm 3 (8 \pm 10)	368.7 \pm 27.6	0.79 \pm 0.08	2.15 \pm 0.22	0.87 \pm 0.19	1,000 (2.1h \pm 29m)
Saturn-GA-ZINC (10)	7 \pm 4 (10 \pm 7)	382.8 \pm 27.9	0.71 \pm 0.08	2.10 \pm 0.15	0.85 \pm 0.17	1,000 (2.0h \pm 26m)
$R_{Double\ MPO}$ (ours)		2 objectives (Docking, AiZynth)				
Saturn-ChEMBL (10)	49 \pm 19 (175 \pm 94)	442.3 \pm 26.2	0.36 \pm 0.05	2.36 \pm 0.17	0.84 \pm 0.06	1,000 (2.0h \pm 31m)
Saturn-GA-ChEMBL (10)	43 \pm 19 (99 \pm 54)	436.1 \pm 17.2	0.39 \pm 0.04	2.35 \pm 0.13	0.77 \pm 0.07	1,000 (1.7h \pm 16m)
Saturn-ZINC (10)	24 \pm 17 (71 \pm 64)	414.0 \pm 20.2	0.52 \pm 0.11	2.30 \pm 0.25	0.90 \pm 0.07	1,000 (1.9h \pm 22m)
Saturn-GA-ZINC (10)	30 \pm 11 (64 \pm 28)	408.1 \pm 12.4	0.46 \pm 0.05	2.19 \pm 0.10	0.86 \pm 0.07	1,000 (1.6h \pm 16m)

197 **discussion on Saturn.** The central message of this section is that molecules satisfying the objective
198 functions can be found within 1,000 oracle calls. An important note is that all RGFN results (top
199 half of Table 1) report results for the top-500 (for Mol. weight, QED⁶², SA score⁴) and top-100 (for
200 AiZynthFinder¹¹⁻¹³) Modes. Saturn does not find 100 Modes in all configurations with 1,000 oracle
201 calls so the metrics are reported for however many Modes were found. Finally, we run Saturn with
202 and without GraphGA-augmented experience replay^{40,68} (see Appendix I for details) and pre-trained
203 with both ChEMBL 33⁶⁴ and ZINC 250k⁶⁵ (see Appendix B for pre-training details). The purpose is
204 to show that the MPO task can be optimized in 1,000 oracle calls using both popular pre-training
205 datasets. We make the following observations: by including AiZynthFinder in the objective function,
206 Saturn generates AiZynthFinder solvable molecules. Including QED and SA score in the objective
207 function also optimizes these metrics (contrast $R_{All\ MPO}$ with $R_{Double\ MPO}$ results). Mol. weight
208 is also implicitly minimized because it is a component of QED. $R_{Double\ MPO}$ finds notably more
209 Modes than $R_{All\ MPO}$ because the optimization task is easier. In all cases, 1/400th the oracle budget
210 is sufficient to find at least *some* molecules that optimize the objectives (and are AiZynthFinder
211 solvable). The wall times are not 1/400th because AiZynthFinder is the slowest oracle, even with
212 multi-threading (see Appendix D). Finally, we highlight that although the raw number of Modes
213 generated when using the $R_{All\ MPO}$ objective function is relatively low (in 1,000 oracle calls), if
214 AiZynthFinder *does* accurately predict "true" synthesizability, then these Modes are immediately
215 actionable. Importantly, they satisfy every metric in the objective function (low docking score, high
216 QED, low SA, and is AiZynthFinder solvable). In practice, one wants to identify a small set of
217 *excellent* candidate molecules as fast as possible (oracle calls and/or wall time). See Appendix G
218 for additional experiments, and particularly how *also* optimizing for QED is a considerably more
219 difficult task.

220 **Qualitative Results.** Fig. 3 shows the docking pose for the generated molecules with the best
221 docking score (no cherry-picking) across all Saturn configurations. In all cases, the pose conforms to
222 the geometry of the binding cavity and the molecule itself is AiZynthFinder solvable (see Appendix
223 F for the solved routes). Generated molecules using $R_{Double\ MPO}$ have better docking scores than




Figure 3: Docked pose of the reference ligand (PDB ID: 7UVU) and generated molecules with the best docking score (DS) across all Saturn configurations and across all 10 seeds (0-9 inclusive). The reference pose is in gray and all generated molecules are in green. **All molecules are AiZynthFinder solvable.** **a.** Molecules generated using $R_{All\ MPO}$. **b.** Molecules generated using $R_{Double\ MPO}$.

224 $R_{All\ MPO}$, which is expected as the optimization task is easier. In the case of $R_{Double\ MPO}$, the
 225 best molecules have docking scores and QED values similar to the best molecules generated by
 226 RGFN²⁹ in 400,000 oracle calls (Fig. 2). We highlight that the molecules from $R_{Double\ MPO}$ possess
 227 extensive carbon rings and are likely exploiting the docking oracle.

228 4.3 Experiment 3: Directly optimizing for synthesizability using AiZynthFinder starting from 229 an unsuitable training distribution

230 Generative models are pre-trained to model the training data distribution. The experiments thus far
 231 use Saturn⁴⁰ which has been trained on either ChEMBL 33⁶⁴ or ZINC 250k⁶⁵. These datasets contain
 232 bio-active molecules and pre-trained models can already generate molecules that can be solved by a
 233 retrosynthesis model². In this section, we pre-train Saturn on the fraction of ZINC 250k that is *not*
 234 AiZynthFinder¹¹⁻¹³ solvable. This model will be referred to as "Purged ZINC" (see Appendix E for
 235 details). The message we convey is that even with an unsuitable training distribution, both $R_{All\ MPO}$
 236 and $R_{Double\ MPO}$ can *still* be optimized under a 1,000 oracle budget.

237 We showcase how curriculum learning (CL)⁷² (decompose a complex optimization objective into
 238 sequential, simpler objectives) can be used by defining two phases of goal-directed generation.
 239 Firstly, "Purged ZINC" is tasked to minimize SA score⁴ (500 oracle budget) as it is correlated with
 240 AiZynthFinder⁴³. Fig. 4a shows the optimization trajectory and the resulting model is referred to as
 241 "Purged ZINC SA". The 500 oracle calls are not counted in the 1,000 oracle budget, as computing
 242 SA score is cheap (this process took 56 seconds). Next, to illustrate distribution learning, we sample
 243 1,000 unique molecules from the "Normal ZINC" (trained on the full dataset), "Purged ZINC",
 244 and "Purged ZINC SA" models and run AiZynthFinder. Fig. 4b shows the fraction of molecules
 245 that are AiZynthFinder solvable. In 56 seconds, the "Purged ZINC" model can be fine-tuned to
 246 immediately generate molecules that are almost all solvable (see Appendix G for additional results
 247 and discussion). We note that "Purged ZINC" still generates molecules that are AiZynthFinder
 248 solvable due to stochastic generation and likely due to the use of SMILES randomization during

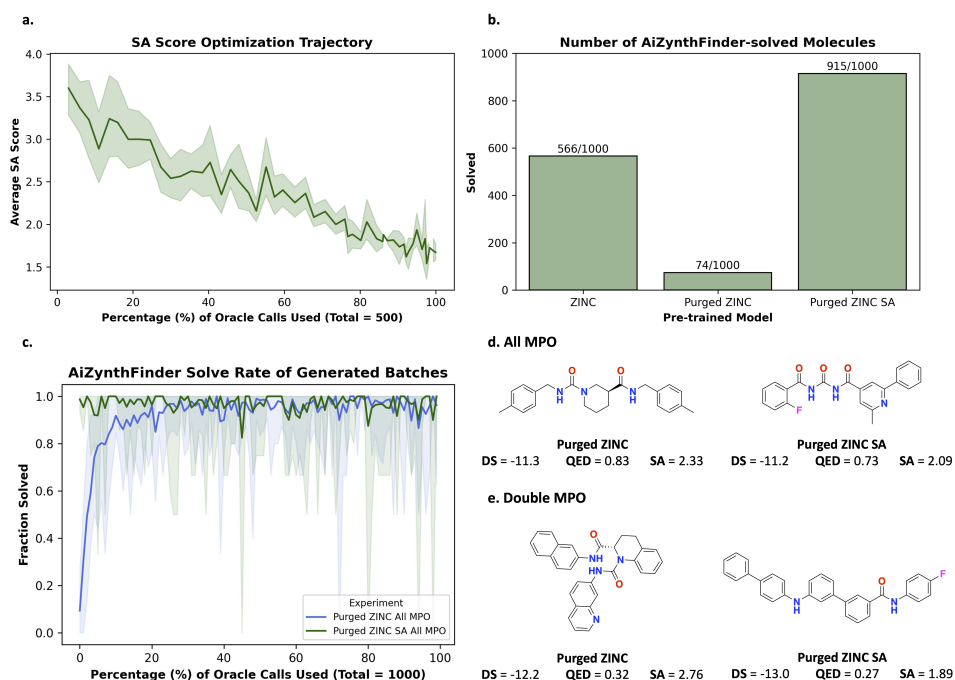


Figure 4: Correlation of SA score and AiZynthFinder solve rate and learning to generate AiZynthFinder solvable molecules. **a.** "Purged ZINC" is tasked to minimize SA score. The average SA score of the sampled batches are shown. **b.** AiZynthFinder solve rates for 1,000 molecules sampled from different models. **c.** All MPO task: fraction of generated molecules (without the GA activated) across all batches that are AiZynthFinder-solvable. Values are the mean and the shaded regions are the minimum-maximum across 10 seeds (0-9 inclusive). **d.** Example molecules generated from the "Purged ZINC" and "Purged ZINC SA" models with the best docking scores (DS).

249 training which enhances chemical space generalizability⁷³. Next, we show how the "Purged ZINC"
 250 model can learn to generate molecules that are AiZynthFinder solvable during the course of RL
 251 (Fig. 4c). We contrast this with the "Purged ZINC SA" model which has almost 100% solve rate
 252 throughout the entire run. During the course of the run, some seeds occasionally generate batches
 253 that are not AiZynthFinder solvable (lower bound of shaded region), but this is not detrimental (see
 254 Appendix J for more details). Fig. 4d shows the molecules with the best docking score generated
 255 across all seeds. The property profiles are essentially the same as the runs with the normal ZINC
 256 model (Fig. 3).

257 **Quantitative Results.** Table 2 contrasts the results of the ClpP docking case study run across 10 seeds
 258 (0-9 inclusive) using the "Normal ZINC", "Purged ZINC", and "Purged ZINC SA" models. Despite
 259 an unsuitable training distribution, "Purged ZINC" can still generate Modes that are AiZynthFinder
 260 solvable, although the solve rate is slightly lower than "Normal ZINC". "Purged ZINC SA" was first
 261 fine-tuned to minimize SA score and already generated mostly AiZynthFinder solvable molecules
 262 (Fig. 4b). This process benefits both $R_{All\ MPO}$ (less so) and $R_{Double\ MPO}$ as the Yield and Modes
 263 found are higher. Next, we highlight that "Purged ZINC" and "Purged ZINC SA" wall times are
 264 longer. This is due to two reasons: firstly, we ran four experiments simultaneously on a single
 265 workstation, which shares resources but makes the total wall time to finish all the experiments faster.
 266 Secondly, "Purged ZINC SA" experiments took longer because the initial CL fine-tuning biases
 267 the model to generate more repeat molecules due to Saturn's⁴⁰ mechanism of local chemical space
 268 exploration. The effect is that it takes longer to exhaust the 1,000 *unique* oracle calls budget.

269 Overall, the property profiles of generated Modes are better than GraphGA⁶⁸, SyntheMol²⁷, FGFN⁶⁹,
 270 and RGFN²⁹. Regardless of the starting model, both $R_{All\ MPO}$ and $R_{Double\ MPO}$ can be optimized
 271 within 1,000 oracle calls. Whether or not the output of AiZynthFinder represents "true" synthesiz-
 272 ability (and the quality of the routes) is beyond the scope of this work. The message we convey
 273 in this section is that generating molecules that are solvable by a retrosynthesis model does not

Table 2: Synthesizability metrics for "Normal ZINC" (results from Table 1), "Purged ZINC", and "Purged ZINC SA". All experiments were run across 10 seeds (0-9 inclusive). The mean and standard deviation are reported. Both Yield and Modes are reported. The number after the configuration denotes the number of successful replicates out of 10 (Modes ≥ 1). The metrics are reported for however many Modes were found.

Method	Modes (Yield)	Mol. weight (\downarrow)	QED (\uparrow)	SA score (\downarrow)	AiZynth (\uparrow)	Oracle calls (Wall time)
R_{AllMPO}		4 objectives (Docking, QED, SA, AiZynth)				
Normal ZINC						
Saturn (9)	6 \pm 3 (8 \pm 10)	368.7 \pm 27.6	0.79 \pm 0.08	2.15 \pm 0.22	0.87 \pm 0.19	1,000 (2.1h \pm 29m)
Saturn-GA (10)	7 \pm 4 (10 \pm 7)	382.8 \pm 27.9	0.71 \pm 0.08	2.10 \pm 0.15	0.85 \pm 0.17	1,000 (2.0h \pm 26m)
Purged ZINC						
Saturn (8)	5 \pm 5 (9 \pm 15)	354.4 \pm 26.2	0.72 \pm 0.15	1.99 \pm 0.27	0.97 \pm 0.05	1,000 (2.8h \pm 72m)
Saturn-GA(10)	10 \pm 3 (14 \pm 5)	381.4 \pm 15.6	0.68 \pm 0.09	2.22 \pm 0.24	0.77 \pm 0.12	1,000 (2.4h \pm 50m)
Purged ZINC SA						
Saturn (10)	9 \pm 5 (16 \pm 11)	365.9 \pm 12.5	0.68 \pm 0.09	1.97 \pm 0.19	0.96 \pm 0.10	1,000 (4.9h \pm 54m)
Saturn-GA (10)	12 \pm 6 (21 \pm 14)	369.7 \pm 15.0	0.69 \pm 0.08	2.06 \pm 0.15	0.89 \pm 0.08	1,000 (3.2h \pm 26m)
$R_{DoubleMPO}$		2 objectives (Docking, AiZynth)				
Normal ZINC						
Saturn (10)	24 \pm 17 (71 \pm 64)	414.0 \pm 20.2	0.52 \pm 0.11	2.30 \pm 0.25	0.90 \pm 0.07	1,000 (1.9h \pm 22m)
Saturn-GA (10)	30 \pm 11 (64 \pm 28)	408.1 \pm 12.4	0.46 \pm 0.05	2.19 \pm 0.10	0.86 \pm 0.07	1,000 (1.6h \pm 16m)
Purged ZINC						
Saturn (10)	27 \pm 19 (114 \pm 107)	425.7 \pm 58.5	0.50 \pm 0.15	2.66 \pm 0.56	0.83 \pm 0.13	1,000 (2.4h \pm 35m)
Saturn-GA (10)	34 \pm 17 (78 \pm 57)	410.8 \pm 16.3	0.44 \pm 0.08	2.29 \pm 0.16	0.78 \pm 0.08	1,000 (2.2h \pm 34m)
Purged ZINC SA						
Saturn (10)	46 \pm 14 (268 \pm 88)	443.5 \pm 31.4	0.39 \pm 0.10	2.13 \pm 0.12	0.87 \pm 0.08	1,000 (3.9h \pm 56m)
Saturn-GA (10)	49 \pm 10 (187 \pm 55)	419.5 \pm 10.6	0.42 \pm 0.04	2.11 \pm 0.05	0.77 \pm 0.06	1,000 (2.9h \pm 29m)

274 *require* synthesizability-constrained design principles. Lastly, we explore the effect of increasing
 275 the oracle budget and implications of heuristics-driven synthesizability and post-hoc retrosynthesis
 276 model filtering in Appendix G.

277 5 Conclusion

278 In this work, we adapt Saturn⁴⁰ which is a sample-efficient autoregressive molecular generative
 279 model using the Mamba⁴¹ architecture to directly optimize for synthesizability using retrosynthesis
 280 models⁴². Our approach contrasts existing works in the field that tackle synthesizability in one
 281 of three ways: goal-directed generation with synthesizability heuristic scores such as SA score⁴,
 282 post-hoc filtering generated molecules with a retrosynthesis model⁶³, or by enforcing synthesizability
 283 design principles in the generative process itself (synthesizability-constrained generation)^{26,28,29}. We
 284 show that with a sufficiently sample-efficient model, treating retrosynthesis models as an oracle is
 285 feasible, and generated molecules can satisfy multi-parameter optimization objectives while being
 286 synthesizable (as deemed by a retrosynthesis model). The main comparison results we show are
 287 on a molecular docking case study proposed by the recent Reaction-GFlowNet (RGFN)²⁹ work
 288 which is a synthesizability-constrained generative model. With 1/400th the oracle budget, Saturn
 289 can generate molecules with better property profiles than GraphGA⁶⁸, SyntheMol²⁷, Fragment
 290 GFlowNet (FGFN)⁶⁹, and RGFN. Moreover, we conduct an artificial experiment to intentionally
 291 purge a training dataset of all molecules that are solvable by the AiZynthFinder¹¹⁻¹³ retrosynthesis
 292 model and pre-train a new model with this dataset. Generated molecules from this model are mostly
 293 not AiZynthFinder solvable, as expected (Fig. 4b). Despite this, we show that within 1/400th the
 294 oracle budget, this model can *still* generate molecules with property profiles better than all comparing
 295 models and are synthesizable (as deemed by AiZynthFinder). The take-home message is that with a
 296 sufficiently sample-efficient model, it is straightforward to treat retrosynthesis models as an oracle
 297 in goal-directed generation. Generating molecules deemed synthesizable by such models does not
 298 *require* synthesizability-constrained generation, which is currently, often *sample-inefficient*.

299 References

- 300 1. Yuanqi Du, Arian R Jamasb, Jeff Guo, Tianfan Fu, Charles Harris, Yingheng Wang, Chenru
301 Duan, Pietro Liò, Philippe Schwaller, and Tom L Blundell. Machine learning-aided generative
302 molecular design. *Nature Machine Intelligence*, pages 1–16, 2024.
- 303 2. Wenhao Gao and Connor W Coley. The synthesizability of molecules proposed by generative
304 models. *Journal of chemical information and modeling*, 60(12):5714–5723, 2020.
- 305 3. Megan Stanley and Marwin Segler. Fake it until you make it? generative de novo design and
306 virtual screening of synthesizable molecules. *Current Opinion in Structural Biology*, 82:102658,
307 2023.
- 308 4. Peter Ertl and Ansgar Schuffenhauer. Estimation of synthetic accessibility score of drug-like
309 molecules based on molecular complexity and fragment contributions. *Journal of cheminforma-
310 tics*, 1:1–11, 2009.
- 311 5. Milan Voršilák, Michal Kolář, Ivan Čmelo, and Daniel Svozil. Syba: Bayesian estimation of
312 synthetic accessibility of organic compounds. *Journal of cheminformatics*, 12:1–13, 2020.
- 313 6. Connor W Coley, Luke Rogers, William H Green, and Klavs F Jensen. Sscore: synthetic
314 complexity learned from a reaction corpus. *Journal of chemical information and modeling*, 58
315 (2):252–261, 2018.
- 316 7. Bowen Liu, Bharath Ramsundar, Prasad Kawthekar, Jade Shi, Joseph Gomes, Quang Luu Nguyen,
317 Stephen Ho, Jack Sloane, Paul Wender, and Vijay Pande. Retrosynthetic reaction prediction
318 using neural sequence-to-sequence models. *ACS central science*, 3(10):1103–1113, 2017.
- 319 8. Marwin HS Segler and Mark P Waller. Neural-symbolic machine learning for retrosynthesis and
320 reaction prediction. *Chemistry—A European Journal*, 23(25):5966–5971, 2017.
- 321 9. Connor W Coley, Luke Rogers, William H Green, and Klavs F Jensen. Computer-assisted
322 retrosynthesis based on molecular similarity. *ACS central science*, 3(12):1237–1245, 2017.
- 323 10. Marwin HS Segler, Mike Preuss, and Mark P Waller. Planning chemical syntheses with deep
324 neural networks and symbolic ai. *Nature*, 555(7698):604–610, 2018.
- 325 11. Amol Thakkar, Thierry Kogej, Jean-Louis Reymond, Ola Engkvist, and Esben Jannik Bjerrum.
326 Datasets and their influence on the development of computer assisted synthesis planning tools in
327 the pharmaceutical domain. *Chemical science*, 11(1):154–168, 2020.
- 328 12. Samuel Genheden, Amol Thakkar, Veronika Chadimová, Jean-Louis Reymond, Ola Engkvist,
329 and Esben Bjerrum. Aizynthfinder: a fast, robust and flexible open-source software for retrosyn-
330 thetic planning. *Journal of cheminformatics*, 12(1):70, 2020.
- 331 13. Lakshidaa Saigiridharan, Alan Kai Hassen, Helen Lai, Paula Torren-Peraire, Ola Engkvist, and
332 Samuel Genheden. Aizynthfinder 4.0: developments based on learnings from 3 years of industrial
333 application. *Journal of Cheminformatics*, 16(1):57, 2024.
- 334 14. Amol Thakkar, Veronika Chadimová, Esben Jannik Bjerrum, Ola Engkvist, and Jean-Louis
335 Reymond. Retrosynthetic accessibility score (rascore)—rapid machine learned synthesizability
336 classification from ai driven retrosynthetic planning. *Chemical science*, 12(9):3339–3349, 2021.
- 337 15. Connor W Coley, Dale A Thomas III, Justin AM Lummiss, Jonathan N Jaworski, Christopher P
338 Breen, Victor Schultz, Travis Hart, Joshua S Fishman, Luke Rogers, Hanyu Gao, et al. A robotic
339 platform for flow synthesis of organic compounds informed by ai planning. *Science*, 365(6453):
340 eaax1566, 2019.
- 341 16. IBM. Rxn for chemistry.
- 342 17. Cheng-Hao Liu, Maksym Korablyov, Stanisław Jastrzebski, Paweł Włodarczyk-Pruszynski,
343 Yoshua Bengio, and Marwin Segler. Retrognn: fast estimation of synthesizability for virtual
344 screening and de novo design by learning from slow retrosynthesis software. *Journal of Chemical
345 Information and Modeling*, 62(10):2293–2300, 2022.

- 346 18. Kevin Yu, Jihye Roh, Ziang Li, Wenhao Gao, Runzhong Wang, and Connor W Coley.
347 Double-ended synthesis planning with goal-constrained bidirectional search. *arXiv preprint*
348 *arXiv:2407.06334*, 2024.
- 349 19. H Maarten Vinkers, Marc R de Jonge, Frederik FD Daeyaert, Jan Heeres, Lucien MH Koymans,
350 Joop H van Lenthe, Paul J Lewi, Henk Timmerman, Koen Van Aken, and Paul AJ Janssen.
351 Synopsis: synthesize and optimize system in silico. *Journal of medicinal chemistry*, 46(13):
352 2765–2773, 2003.
- 353 20. Markus Hartenfeller, Heiko Zettl, Miriam Walter, Matthias Rupp, Felix Reisen, Ewgenij
354 Proschak, Sascha Weggen, Holger Stark, and Gisbert Schneider. Dogs: reaction-driven de
355 novo design of bioactive compounds. *PLoS computational biology*, 8(2):e1002380, 2012.
- 356 21. John Bradshaw, Brooks Paige, Matt J Kusner, Marwin Segler, and José Miguel Hernández-Lobato.
357 A model to search for synthesizable molecules. *Advances in Neural Information Processing*
358 *Systems*, 32, 2019.
- 359 22. John Bradshaw, Brooks Paige, Matt J Kusner, Marwin Segler, and José Miguel Hernández-
360 Lobato. Barking up the right tree: an approach to search over molecule synthesis dags. *Advances*
361 *in neural information processing systems*, 33:6852–6866, 2020.
- 362 23. Ksenia Korovina, Sailun Xu, Kirthevasan Kandasamy, Willie Neiswanger, Barnabas Poczos,
363 Jeff Schneider, and Eric Xing. Chembo: Bayesian optimization of small organic molecules
364 with synthesizable recommendations. In *International Conference on Artificial Intelligence and*
365 *Statistics*, pages 3393–3403. PMLR, 2020.
- 366 24. Sai Krishna Gottipati, Boris Sattarov, Sufeng Niu, Yashaswi Pathak, Haoran Wei, Shengchao
367 Liu, Simon Blackburn, Karam Thomas, Connor Coley, Jian Tang, et al. Learning to navigate the
368 synthetically accessible chemical space using reinforcement learning. In *International conference*
369 *on machine learning*, pages 3668–3679. PMLR, 2020.
- 370 25. Julien Horwood and Emmanuel Noutahi. Molecular design in synthetically accessible chemical
371 space via deep reinforcement learning. *ACS omega*, 5(51):32984–32994, 2020.
- 372 26. Wenhao Gao, Rocío Mercado, and Connor W Coley. Amortized tree generation for bottom-up
373 synthesis planning and synthesizable molecular design. *Proc. 10th International Conference on*
374 *Learning Representations*, 2022.
- 375 27. Kyle Swanson, Gary Liu, Denise B Catacutan, Autumn Arnold, James Zou, and Jonathan M
376 Stokes. Generative ai for designing and validating easily synthesizable and structurally novel
377 antibiotics. *Nature Machine Intelligence*, 6(3):338–353, 2024.
- 378 28. Miruna Cretu, Charles Harris, Julien Roy, Emmanuel Bengio, and Pietro Liò. Synflownet:
379 Towards molecule design with guaranteed synthesis pathways. *arXiv preprint arXiv:2405.01155*,
380 2024.
- 381 29. Michał Koziarski, Andrei Rekes, Dmytro Shevchuk, Almer van der Sloot, Piotr Gaiński, Yoshua
382 Bengio, Cheng-Hao Liu, Mike Tyers, and Robert A Batey. Rgfn: Synthesizable molecular
383 generation using gflownets. *arXiv preprint arXiv:2406.08506*, 2024.
- 384 30. Wenhao Gao, Tianfan Fu, Jimeng Sun, and Connor Coley. Sample efficiency matters: a
385 benchmark for practical molecular optimization. *Advances in neural information processing*
386 *systems*, 35:21342–21357, 2022.
- 387 31. Soojung Yang, Doyeong Hwang, Seul Lee, Seongok Ryu, and Sung Ju Hwang. Hit and lead
388 discovery with explorative rl and fragment-based molecule generation. *Advances in Neural*
389 *Information Processing Systems*, 34:7924–7936, 2021.
- 390 32. Tianfan Fu, Wenhao Gao, Connor Coley, and Jimeng Sun. Reinforced genetic algorithm for
391 structure-based drug design. *Advances in Neural Information Processing Systems*, 35:12325–
392 12338, 2022.

- 393 33. Seul Lee, Jaehyeong Jo, and Sung Ju Hwang. Exploring chemical space with score-based out-of-
394 distribution generation. In *International Conference on Machine Learning*, pages 18872–18892.
395 PMLR, 2023.
- 396 34. Seul Lee, Seanie Lee, and Sung Ju Hwang. Drug discovery with dynamic goal-aware fragments.
397 *arXiv preprint arXiv:2310.00841*, 2023.
- 398 35. Tony Shen, Mohit Pandey, and Martin Ester. Tacogfn: Target conditioned gflownet for drug
399 design. In *NeurIPS 2023 Generative AI and Biology (GenBio) Workshop*, 2023.
- 400 36. Jeff Guo, Franziska Knuth, Christian Margreitter, Jon Paul Janet, Kostas Papadopoulos, Ola
401 Engkvist, and Atanas Patronov. Link-invent: generative linker design with reinforcement learning.
402 *Digital Discovery*, 2(2):392–408, 2023.
- 403 37. Michael Dodds, Jeff Guo, Thomas Löhr, Alessandro Tibo, Ola Engkvist, and Jon Paul Janet.
404 Sample efficient reinforcement learning with active learning for molecular design. *Chemical*
405 *Science*, 15(11):4146–4160, 2024.
- 406 38. Jeff Guo and Philippe Schwaller. Augmented memory: Sample-efficient generative molecular
407 design with reinforcement learning. *JACS Au*, 2024.
- 408 39. Jeff Guo and Philippe Schwaller. Beam enumeration: Probabilistic explainability for sample
409 efficient self-conditioned molecular design. In *Proc. 12th International Conference on Learning*
410 *Representations*, 2024.
- 411 40. Jeff Guo and Philippe Schwaller. Saturn: Sample-efficient generative molecular design using
412 memory manipulation. *arXiv preprint arXiv:2405.17066*, 2024.
- 413 41. Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces.
414 *arXiv preprint arXiv:2312.00752*, 2023.
- 415 42. Albin Ekborg. De novo molecular generation of molecules with consistent synthetic strategy.
416 Master’s thesis, Chalmers University of Technology, 2024.
- 417 43. Grzegorz Skoraczyński, Mateusz Kitlas, Błażej Miasojedow, and Anna Gambin. Critical as-
418 sessment of synthetic accessibility scores in computer-assisted synthesis planning. *Journal of*
419 *Cheminformatics*, 15(1):6, 2023.
- 420 44. Rebecca M Neeser, Bruno Correia, and Philippe Schwaller. Fsscore: A machine learning-based
421 synthetic feasibility score leveraging human expertise. *arXiv preprint arXiv:2312.12737*, 2023.
- 422 45. Oh-Hyeon Choung, Riccardo Vianello, Marwin Segler, Nikolaus Stiefl, and José Jiménez-Luna.
423 Extracting medicinal chemistry intuition via preference machine learning. *Nature Communica-*
424 *tions*, 14(1):6651, 2023.
- 425 46. Sara Szymkuć, Ewa P Gajewska, Tomasz Klucznik, Karol Molga, Piotr Dittwald, Michał Startek,
426 Michał Bajczyk, and Bartosz A Grzybowski. Computer-assisted synthetic planning: the end of
427 the beginning. *Angewandte Chemie International Edition*, 55(20):5904–5937, 2016.
- 428 47. Bartosz A Grzybowski, Sara Szymkuć, Ewa P Gajewska, Karol Molga, Piotr Dittwald, Agnieszka
429 Wołos, and Tomasz Klucznik. Chematica: a story of computer code that started to think like a
430 chemist. *Chem*, 4(3):390–398, 2018.
- 431 48. Ian A Watson, Jibo Wang, and Christos A Nicolaou. A retrosynthetic analysis algorithm
432 implementation. *Journal of cheminformatics*, 11:1–12, 2019.
- 433 49. Molecule.one. The m1 platform.
- 434 50. Philippe Schwaller, Riccardo Petraglia, Valerio Zullo, Vishnu H Nair, Rico Andreas Haeuselmann,
435 Riccardo Pisoni, Costas Bekas, Anna Iuliano, and Teodoro Laino. Predicting retrosynthetic
436 pathways using transformer-based models and a hyper-graph exploration strategy. *Chemical*
437 *science*, 11(12):3316–3325, 2020.

- 438 51. Amol Thakkar, Alain C Vaucher, Andrea Byekwaso, Philippe Schwaller, Alessandra Toniato,
439 and Teodoro Laino. Unbiasing retrosynthesis language models with disconnection prompts. *ACS*
440 *Central Science*, 9(7):1488–1498, 2023.
- 441 52. Gian Marco Ghiandoni, Michael J Bodkin, Beining Chen, Dimitar Hristozov, James EA Wallace,
442 James Webster, and Valerie J Gillet. Renate: a pseudo-retrosynthetic tool for synthetically
443 accessible de novo design. *Molecular Informatics*, 41(4):2100207, 2022.
- 444 53. Vendy Fialková, Jiayi Zhao, Kostas Papadopoulos, Ola Engkvist, Esben Jannik Bjerrum, Thierry
445 Kogej, and Atanas Patronov. Libinvent: reaction-based generative scaffold decoration for in
446 silico library design. *Journal of Chemical Information and Modeling*, 62(9):2046–2063, 2021.
- 447 54. Yoshua Bengio, Salem Lahlou, Tristan Deleu, Edward J Hu, Mo Tiwari, and Emmanuel Bengio.
448 Gflownet foundations. *Journal of Machine Learning Research*, 24(210):1–55, 2023.
- 449 55. Shitong Luo, Wenhao Gao, Zuofan Wu, Jian Peng, Connor W Coley, and Jianzhu Ma. Projecting
450 molecules into synthesizable chemical spaces. *Proc. 41st International Conference on Machine*
451 *Learning*, 2024.
- 452 56. Karol Molga, Ewa P Gajewska, Sara Szymkuć, and Bartosz A Grzybowski. The logic of
453 translating chemical knowledge into machine-processable forms: a modern playground for
454 physical-organic chemistry. *Reaction Chemistry & Engineering*, 4(9):1506–1521, 2019.
- 455 57. Barbara Mikulak-Klucznik, Patrycja Gołębiowska, Alison A Bayly, Oskar Popik, Tomasz
456 Klucznik, Sara Szymkuć, Ewa P Gajewska, Piotr Dittwald, Olga Staszewska-Krajewska, Wiktor
457 Beker, et al. Computational planning of the synthesis of complex natural products. *Nature*, 588
458 (7836):83–88, 2020.
- 459 58. Oleksandr O Grygorenko, Dmytro S Radchenko, Igor Dziuba, Alexander Chuprina, Kateryna E
460 Gubina, and Yurii S Moroz. Generating multibillion chemical space of readily accessible
461 screening compounds. *Iscience*, 23(11), 2020.
- 462 59. Oleg Trott and Arthur J Olson. Autodock vina: improving the speed and accuracy of docking
463 with a new scoring function, efficient optimization, and multithreading. *Journal of computational*
464 *chemistry*, 31(2):455–461, 2010.
- 465 60. Amr Alhossary, Stephanus Daniel Handoko, Yuguang Mu, and Chee-Keong Kwoh. Fast, accurate,
466 and reliable molecular docking with quickvina 2. *Bioinformatics*, 31(13):2214–2216, 2015.
- 467 61. Shidi Tang, Ji Ding, Xiangyu Zhu, Zheng Wang, Haitao Zhao, and Jiansheng Wu. Vina-gpu
468 2.1: towards further optimizing docking speed and precision of autodock vina and its derivatives.
469 *bioRxiv*, pages 2023–11, 2023.
- 470 62. G Richard Bickerton, Gaia V Paolini, Jérémy Besnard, Sorel Muresan, and Andrew L Hopkins.
471 Quantifying the chemical beauty of drugs. *Nature chemistry*, 4(2):90–98, 2012.
- 472 63. Jason D Shields, Rachel Howells, Gillian Lamont, Yin Leilei, Andrew Madin, Christopher E
473 Reimann, Hadi Rezaei, Tristan Reuillon, Bryony Smith, Clare Thomson, et al. Aizynth impact
474 on medicinal chemistry practice at astrazeneca. *RSC Medicinal Chemistry*, 15(4):1085–1095,
475 2024.
- 476 64. Anna Gaulton, Louisa J Bellis, A Patricia Bento, Jon Chambers, Mark Davies, Anne Hersey,
477 Yvonne Light, Shaun McGlinchey, David Michalovich, Bissan Al-Lazikani, et al. ChEMBL:
478 a large-scale bioactivity database for drug discovery. *Nucleic acids research*, 40(D1):D1100–
479 D1107, 2012.
- 480 65. Teague Sterling and John J Irwin. Zinc 15–ligand discovery for everyone. *Journal of chemical*
481 *information and modeling*, 55(11):2324–2337, 2015.
- 482 66. Martin Buttenschoen, Garrett M Morris, and Charlotte M Deane. Posebusters: Ai-based docking
483 methods fail to generate physically valid poses or generalise to novel sequences. *Chemical*
484 *Science*, 15(9):3130–3139, 2024.

- 485 67. John J Irwin, Da Duan, Hayarpi Torosyan, Allison K Doak, Kristin T Ziebart, Teague Sterling,
486 Gurgen Tumanian, and Brian K Shoichet. An aggregation advisor for ligand discovery. *Journal*
487 *of medicinal chemistry*, 58(17):7076–7087, 2015.
- 488 68. Jan H Jensen. A graph-based genetic algorithm and generative model/monte carlo tree search for
489 the exploration of chemical space. *Chemical science*, 10(12):3567–3572, 2019.
- 490 69. Emmanuel Bengio, Moksh Jain, Maksym Korablyov, Doina Precup, and Yoshua Bengio. Flow
491 network based generative models for non-iterative diverse candidate generation. *Advances in*
492 *Neural Information Processing Systems*, 34:27381–27394, 2021.
- 493 70. John A Arnott and Sonia Lobo Planey. The influence of lipophilicity in drug discovery and
494 design. *Expert opinion on drug discovery*, 7(10):863–875, 2012.
- 495 71. Jeff Guo, Jon Paul Janet, Matthias R Bauer, Eva Nittinger, Kathryn A Giblin, Kostas Papadopou-
496 los, Alexey Voronov, Atanas Patronov, Ola Engkvist, and Christian Margreitter. Dockstream: a
497 docking wrapper to enhance de novo molecular design. *Journal of cheminformatics*, 13:1–21,
498 2021.
- 499 72. Jeff Guo, Vendy Fialková, Juan Diego Arango, Christian Margreitter, Jon Paul Janet, Kostas
500 Papadopoulos, Ola Engkvist, and Atanas Patronov. Improving de novo molecular design with
501 curriculum learning. *Nature Machine Intelligence*, 4(6):555–563, 2022.
- 502 73. Josep Arús-Pous, Simon Viet Johansson, Oleksii Prykhodko, Esben Jannik Bjerrum, Christian
503 Tyrchan, Jean-Louis Reymond, Hongming Chen, and Ola Engkvist. Randomized smiles strings
504 improve the quality of molecular generative models. *Journal of cheminformatics*, 11:1–13, 2019.
- 505 74. Esben Jannik Bjerrum. Smiles enumeration as data augmentation for neural network modeling
506 of molecules. *arXiv preprint arXiv:1703.07076*, 2017.
- 507 75. Peter Eastman, Mark S Friedrichs, John D Chodera, Randall J Radmer, Christopher M Bruns,
508 Joy P Ku, Kyle A Beauchamp, Thomas J Lane, Lee-Ping Wang, Diwakar Shukla, et al. Openmm
509 4: a reusable, extensible, hardware independent library for high performance molecular simula-
510 tion. *Journal of chemical theory and computation*, 9(1):461–469, 2013.
- 511 76. Sereina Riniker and Gregory A Landrum. Better informed distance geometry: using what we
512 know to improve conformation generation. *Journal of chemical information and modeling*, 55
513 (12):2562–2574, 2015.
- 514 77. Anthony K Rappé, Carla J Casewit, KS Colwell, William A Goddard III, and W Mason Skiff.
515 Uff, a full periodic table force field for molecular mechanics and molecular dynamics simulations.
516 *Journal of the American chemical society*, 114(25):10024–10035, 1992.
- 517 78. Noel M O’Boyle, Michael Banck, Craig A James, Chris Morley, Tim Vandermeersch, and
518 Geoffrey R Hutchison. Open babel: An open chemical toolbox. *Journal of cheminformatics*, 3:
519 1–14, 2011.
- 520 79. Binghong Chen, Chengtao Li, Hanjun Dai, and Le Song. Retro*: learning retrosynthetic planning
521 with neural guided a* search. In *International conference on machine learning*, pages 1608–1616.
522 PMLR, 2020.
- 523 80. Krzysztof Maziarz, Austin Tripp, Guoqing Liu, Megan Stanley, Shufang Xie, Piotr Gaiński,
524 Philipp Seidl, and Marwin Segler. Re-evaluating retrosynthesis algorithms with synthesus.
525 *arXiv preprint arXiv:2310.19796*, 2023.
- 526 81. Alan Kai Hassen, Martin Sicho, Yorick J van Aalst, Mirjam CW Huizenga, Darcy NR Reynolds,
527 Sohvi Luukkonen, Andrius Bernatavicius, Djork-Arné Clevert, Antonius PA Janssen, Gerard JP
528 van Westen, et al. Generate what you can make: Achieving in-house synthesizability with readily
529 available resources in de novo drug design. *ChemRxiv*, 2024.
- 530 82. David Weininger. Smiles, a chemical language and information system. 1. introduction to
531 methodology and encoding rules. *Journal of chemical information and computer sciences*, 28(1):
532 31–36, 1988.

533 Appendix

534 The Appendix contains details on the procedure we took to reproduce RGFN's²⁹ oracle as the code is
535 not released. In addition, we report the computational resources used, how Saturn was pre-trained,
536 and AiZynthFinder execution details.

537 A Compute Resources

538 All experiments were run on a single workstation with an NVIDIA RTX A6000 GPU 48GB memory
539 and AMD Ryzen 9 5900X 24-Core CPU. 48GB GPU memory is not required. QuickVina2-GPU-
540 2.1⁵⁹⁻⁶¹ with 'thread' = 8,000 (following the RGFN²⁹ work) takes up to 12GB GPU memory. We
541 further note that Saturn's wall times reported in Table 1 are longer than actually required as we always
542 run 2-4 experiments in parallel, which share the workstation's resources, but makes the *total* wall
543 time less.

544 B Saturn Pre-training Details

545 This section contains the exact protocol used for Saturn pre-training on ChEMBL 33⁶⁴ and ZINC
546 250k⁶⁵. The details and pre-trained models are taken from the original Saturn⁴⁰ paper and included
547 here.

548 B.1 ChEMBL 33

549 Each step is followed by the SMILES remaining after the filtering step.

- 550 1. Download raw ChEMBL 33 - 2,372,674
- 551 2. Standardization (charge and isotope handling) based on [https://github.com/
552 MolecularAI/ReinventCommunity/blob/master/notebooks/Data_Preparation.
553 ipynb](https://github.com/MolecularAI/ReinventCommunity/blob/master/notebooks/Data_Preparation.ipynb). All SMILES that could not be parsed by RDKit were removed - 2,312,459
- 554 3. Kept only the unique SMILES - 2,203,884
- 555 4. Tokenize all SMILES based on REINVENT's tokenizer: [https://github.com/
556 MolecularAI/reinvent-models/blob/main/reinvent_models/reinvent_core/
557 models/vocabulary.py](https://github.com/MolecularAI/reinvent-models/blob/main/reinvent_models/reinvent_core/models/vocabulary.py)
- 558 5. Keep SMILES ≤ 80 tokens - 2,065,099
- 559 6. $150 \leq$ molecular weight ≤ 600 - 2,016,970
- 560 7. Number of heavy atoms ≤ 40 - 1,975,282
- 561 8. Number of rings ≤ 8 - 1,974,522
- 562 9. Size of largest ring ≤ 8 - 1,961,690
- 563 10. Longest aliphatic carbon chain ≤ 5 - 1,950,213
- 564 11. Removed SMILES containing the following tokens (due to undesired chemistry and low
565 token frequency): [S+], [C-], [s+], [O], [S@+], [S@@+], [S-], [o+], [NH+], [n-], [N@],
566 [N@@], [N@+], [N@@+], [S@@], [C+], [S@], [c+], [NH2+], [SH], [NH-], [cH-], [O+],
567 [c-], [CH], [SH+], [CH2-], [OH+], [nH+], [SH2] - **1,942,081**

568 The final vocabulary contained 37 tokens (2 extra tokens were added, indicating <START> and
569 <END>).

570 The Mamba model has 5,265,920 parameters. The hyperparameters are the default parameters in the
571 code base.

572 **The pre-training parameters were:**

- 573 1. Max training steps = 20 (each training step entails a full pass through the dataset)
- 574 2. Seed = 0
- 575 3. Batch size = 512

- 576 4. Learning rate = 0.0001
577 5. Randomize⁷⁴ every batch of SMILES

578 The following checkpoint was used: Epoch 18, NLL = 32.21, Validity (10k) = 95.60%.

579 B.2 ZINC 250k

580 ZINC 250k⁶⁵ was downloaded and used as is.

581 **The pre-training parameters were:**

- 582 1. Training steps = 50 (each training step entails a full pass through the dataset)
583 2. Seed = 0
584 3. Batch size = 512
585 4. Learning rate = 0.0001
586 5. Train with SMILES randomization⁷⁴ (all SMILES in each batch was randomized)

587 The final vocabulary contained 66 tokens (2 extra tokens were added, indicating <START> and
588 <END>).

589 The Mamba model has 5,272,832 parameters (slightly larger than ChEMBL 33 model because the
590 vocabulary size here is larger). The following checkpoint was used: Epoch 50, NLL = 28.10, Validity
591 (10k) = 95.20%.

592 C Reproducing RGFN's Oracle

593 This section contains the steps we took to reproduce RGFN's²⁹ ATP-dependent Clp protease prote-
594 olytic subunit (ClpP) docking case study as faithfully as we could.

595 **Target Preparation.** Following Appendix C.1 of the RGFN paper, we downloaded the 7UVU ClpP
596 crystal structure here: <https://www.rcsb.org/structure/7UVU>. All molecules (complexed
597 inhibitors, solvents, etc.) were removed, keeping only two monomeric units. Two structures were
598 saved: The apo protein (no other molecules present) and the reference ligand. **The following step
599 differs from RGFN:** the apo protein was processed with PDBFixer⁷⁵ to fix missing atoms and
600 residues. We performed this step because errors were thrown during docking when using the raw apo
601 protein structure.

602 **Docking Details.** We implement QuickVina2-GPU-2.1⁵⁹⁻⁶¹ following the instructions in the GitHub
603 repository here: <https://github.com/DeltaGroupNJUPT/Vina-GPU-2.1>. The reference ligand
604 structure that was saved out in the previous step is used here to define the docking box. Specifically,
605 the average coordinates of the ligand denote the docking centroid. **The following may differ from
606 RGFN:** We define the docking box as 20 Å x 20 Å x 20 Å as it was unclear how it should be defined
607 based on RGFN's protocol. This box size has worked on many other protein targets⁷¹ when docking
608 with AutoDock Vina⁵⁹ which is the predecessor of QuickVina2-GPU-2.1.

609 **Docking Workflow.** Following RGFN's protocol, QuickVina2-GPU-2.1 used the following param-
610 eters: 'thread' = 8,000 with 'search depth' = "heuristic" which is the default. Next, all ligands were
611 docked following RGFN's workflow:

- 612 1. Start with batch of generated SMILES from Saturn
613 2. Canonicalize the SMILES
614 3. Convert to RDKit Mol objects
615 4. Protonate the Mols
616 5. Generate 1 (lowest energy) conformer using 'ETKDG'⁷⁶
617 6. Minimize energy with the Universal Force Field (UFF)⁷⁷
618 7. Write out the conformers as 'PDB' files
619 8. Using Open Babel⁷⁸, convert the 'PDB' to 'PDBQT' format

620 9. Execute QuickVina2-GPU-2.1 docking

621 **Protocol Validation.** We make further efforts to ensure the oracle is as faithful as possible to RGFN’s
622 implementation. When executing QuickVina2-GPU-2.1, if a seed is not specified, a random seed is
623 used. It is unclear if a seed was set in the RGFN²⁹ work. In our experiments, the seed is 0. We re-dock
624 the reference ligand and find that the pose is similar to Figure 15 in the RGFN work. However, the
625 docking score we obtain is -9.2 whereas the RGFN work reports -10.31. Subsequently, we execute
626 docking 100 times (letting QuickVina2-GPU 2.1 select the random seed) and observed that seed =
627 448029751 gives a similar pose to RGFN’s pose and yields a docking score of -10.1. We additionally
628 found that seed = 1920393356 yields a docking score of -10.3 but the pose is reflected. Finally seed
629 = 673697018 yields a docking score of -8.2 and is a completely different pose. It is intractable to try
630 every seed.

631 Therefore, we end this section by stating that it is hard to say if we *exactly* re-implement RGFN’s²⁹
632 docking oracle. However, we believe it still enables us to convey the primary message of our work:
633 retrosynthesis models can be directly treated as an oracle and be explicitly optimized for during
634 generation.

635 D AiZynthFinder

636 AiZynthFinder¹¹⁻¹³ was used as is, without modification. The source code was cloned from the
637 GitHub repository here: <https://github.com/MolecularAI/aizynthfinder>. The environ-
638 ment and package were installed following the README. Following the documentation here:
639 <https://molecularai.github.io/aizynthfinder/>, we downloaded the public data and used
640 AiZynthFinder as is. Every batch of molecules generated by Saturn (16 at max) is chunked into 4
641 sub-sets for multi-thread execution. Finally, we consider a molecule AiZynthFinder "solvable" if
642 the "is_solved" flag is True. This flag denotes whether the top scored (accounting for tree depth and
643 fraction of building blocks in stock)^{12,13} is solved.

644 E AiZynthFinder purged ZINC 250k Pre-training Details.

645 In Experiment 3, we pre-train Saturn on a sub-set of ZINC 250k⁶⁵ that *is not* AiZynthFinder¹¹⁻¹³
646 solvable. The goal is to show that Saturn can *still* optimize for generating molecules that are
647 AiZynthFinder solvable despite being trained on no molecules that can be.

648 **Purged Dataset.** We first run AiZynthFinder on the entirety of ZINC 250k on a single workstation
649 with an NVIDIA RTX A6000 GPU 48GB memory and AMD Ryzen 9 5900X 24-Core CPU. The
650 process was run using multi-threading across 12 workers and took 62 hours. We save the unique
651 SMILES (98,110) of all the molecules that *are not* AiZynthFinder solvable. This is the dataset used
652 for pre-training.

653 **Pre-training.** Following the same pre-training parameters used in the original Saturn⁴⁰ work:

- 654 1. Training steps = 100 (each training step entails a full pass through the dataset)
- 655 2. Seed = 0
- 656 3. Batch size = 512
- 657 4. Learning rate = 0.0001
- 658 5. Train with SMILES randomization⁷⁴ (all SMILES in each batch was randomized)

659 The final vocabulary contained 57 tokens (2 extra tokens were added, indicating <START> and
660 <END>). This is less than the normal ZINC 250k model (66 tokens) because some tokens are not
661 present in the purged dataset.

662 The Mamba model has 5,271,040 parameters (less than the normal 250k model because the vocabulary
663 size is smaller). The following checkpoint was used: Epoch 100, NLL = 27.78, Validity (10k) =
664 92.27% and the training time was 4.7 hours.




Figure F5: AiZynthFinder solved routes (top-scoring) for All MPO example molecules.




Figure F6: AiZynthFinder solved routes (top-scoring) for Double MPO example molecules.

665 F AiZynthFinder Routes

666 The AiZynthFinder solved routes for the 8 example molecules shown in Fig. 3 are shown here. The
667 All MPO routes (Fig. F5) are generally shorter than the Double MPO routes (Fig. F6). This suggests
668 that enforcing QED and SA score also implicitly makes the predicted forward syntheses shorter.
669 We note that it is possible to design an objective function that also aims to generate short paths by
670 rewarding short paths. We do not explore this here and leave it for future work.

671 G Supplementary Results

672 In this section, supplementary results are reported which aim to address/provide evidence for three
673 points:

- 674 1. Effect of increasing the oracle budget when optimizing AiZynthFinder
- 675 2. *Jointly* optimizing QED with docking score is *considerably* more difficult than just optimiz-
676 ing docking score
- 677 3. Optimizing SA score *can* be a better allocation of computational resources

Table 3: Synthesizability metrics across various Saturn experiments. Metrics are reported for however many Modes are found. For these supplemental results, only one replicate is performed with seed = 0.

Method	Modes (Yield)	Mol. weight (\downarrow)	QED (\uparrow)	SA score (\downarrow)	AiZynth (\uparrow)	Oracle calls (Wall time)
<i>R_{All MPO}</i>		4 objectives (Docking, QED, SA, AiZynth)				
Saturn-GA-ChEMBL	108 (222)	370.9	0.84	2.44	0.70	5,000 (23.6h)
Saturn-GA-ZINC	74 (230)	371.1	0.81	2.45	0.69	5,000 (21.2h)
<i>R_{Double MPO}</i>		2 objectives (Docking, AiZynth)				
Saturn-ChEMBL	302 (3804)	486.2	0.28	2.40	0.82	5,000 (21.4h)
Saturn-GA-ChEMBL	323 (3053)	464.1	0.34	2.51	0.68	5,000 (11.3h)
Saturn-ZINC	266 (2783)	521.2	0.25	2.40	0.76	5,000 (13.5h)
Saturn-GA-ZINC	327 (2741)	455.6	0.34	2.48	0.72	5,000 (11.3h)
<i>R_{All MPO}</i> (but without AiZynth)		3 objectives (Docking, QED, SA)				
Saturn-ChEMBL	332 (1219)	376.7	0.80	2.66	0.39	10,000 (2.4h)
Saturn-ZINC	332 (1108)	382.1	0.76	2.43	0.55	10,000 (2.3h)
<i>R_{RGFN}</i> - Results from Fig. 2		1 objective (Docking)				
Saturn-ChEMBL	469 (8389)	511.9	0.26	3.09	0.14	10,000 (2.2h)

678 **Increasing the oracle budget leads to notably increased wall times.** In the main text results,
679 *R_{All MPO}* does not find that many Modes. We investigate the effect of increasing the oracle budget
680 (Table 3) with the GA activated (which recover diversity so as to satisfy the Modes criterion that
681 Modes must have < 0.5 Tanimoto similarity with other Modes). More Modes are found but the wall
682 time is *drastically* higher. With 5x the oracle budget (5,000 compared to 1,000 in the main text),
683 one may expect 5x the wall time (12-15 hours) but the wall time is almost 24 hours. The reason
684 is due to Saturn’s sampling behaviour which locally explores chemical space⁴⁰. The parameters
685 of Saturn could be changed to loosen this local exploration behaviour but we do not explore this.
686 We demonstrate the application of Saturn out-of-the-box. As a consequence of this, many repeat
687 molecules are generated, which do not impose an oracle call as the reward is retrieved from an
688 oracle cache, but makes the sampled batch (new molecules) smaller. Multi-threading was used to
689 run AiZynthFinder faster (see Appendix D for more details). Consider batches of 1 molecule and
690 4 molecules. This can take a similar wall time as molecules can be chunked, thus benefiting from
691 multi-threading. This could be mitigated, for example, by using a faster retrosynthesis model which
692 can come with advantages and disadvantages^{79,80} and/or CPU parallelization. Finally, we highlight
693 that deactivating the GA will likely lead to higher Yield and AiZynthFinder solve rate, as shown in
694 Tables 1 and 2. We reiterate that activating the GA was to satisfy the Mode metric.

695 **Jointly optimizing QED with docking score is considerably more difficult than just optimizing**
696 **docking score.** RGFN²⁹ reports their mean and standard deviation of QED values as 0.23 ± 0.04

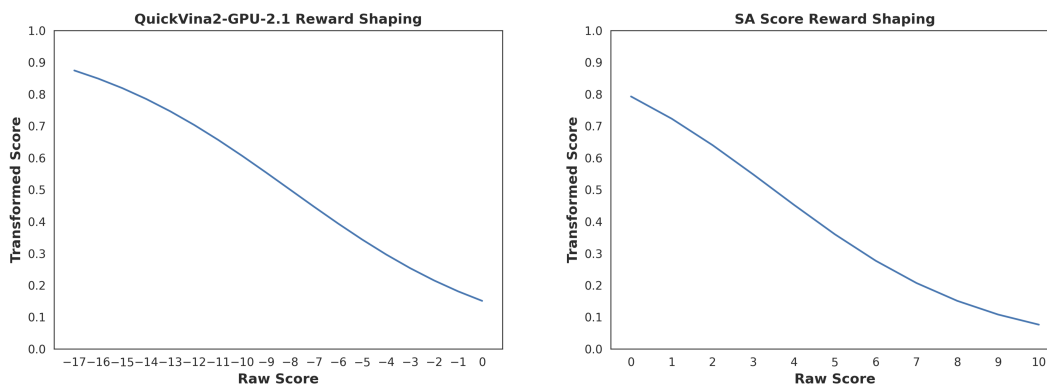


Figure H7: Saturn reward shaping functions for QuickVina2-GPU-2.1 and SA score.

697 (unclear how many replicates this was over). This is low and suggests the model is exploiting
 698 the docking algorithm as shown in Fig. 2. To show that *jointly* optimizing QED and docking is
 699 a *considerably* more difficult task, we first cross-reference the results for $R_{Double\ MPO}$ (Table 3)
 700 where the Modes and Yield are notably higher than $R_{All\ MPO}$. Next, we cross-reference the results
 701 when QED is *not* being optimized (Table 3 last row). The Yield is *much* higher (molecules with
 702 docking score < 10,000) but the QED values are similar to RGFN, which again, suggests the docking
 703 algorithm is being exploited.

704 **Optimizing SA score can be a better allocation of computational resources.** SA score⁴ is
 705 correlated with AiZynthFinder solve rate⁴³. In the main text Fig 4, we empirically demonstrate
 706 this, as 56 seconds of fine-tuning a pre-trained model that has *never* seen an AiZynthFinder solved
 707 molecule, results in a model that generates molecules almost all solvable. The natural next question
 708 is, would simply optimizing SA score be a better allocation of computational resources (as is
 709 commonly done)? Under the same wall time, many more queries to SA score can be made because
 710 it is computationally cheap. Correspondingly, we use the $R_{All\ MPO}$ objective function but omit
 711 AiZynthFinder (only docking, QED, and SA score) and run the ChEMBL and ZINC pre-trained
 712 models for 10,000 oracle calls (Table 3). Firstly, the wall time is similar to running 1,000 oracle
 713 calls of AiZynthFinder (cross-reference Table 1). Next, while a smaller fraction of the Modes are
 714 AiZynthFinder solvable, the raw number is higher than directly optimizing AiZynthFinder. This
 715 reinforces that post-hoc retrosynthesis model filtering is valid and is often what is done in practice⁶³.
 716 Crucially, the actual percentage of AiZynthFinder solve rate may not *actually* matter. What matters is
 717 that a user can reasonably expect a generative model to generate molecules satisfying the objective
 718 function within the allotted oracle budget and/or wall time. In this specific example, it does not matter
 719 that Saturn-ZINC "only" has 55% solve rate when optimizing docking, QED, and SA score (Table 3).
 720 Running the 332 Modes through AiZynthFinder only took about 20 minutes (about 183/332 can be
 721 solved). A user would only care that in under 3 hours, 183 Modes were found that have low docking
 722 score, high QED, low SA score, and are AiZynthFinder solvable.

723 Finally, we wish to be prudent with making definitive statements about whether just optimizing SA
 724 score is strictly *better* than including a retrosynthesis model in the objective function. In this section
 725 alone, we have highlighted that different retrosynthesis models can have a large impact on wall
 726 time^{79,80}, where faster wall times would narrow the gap between SA score's wall time. Moreover,
 727 molecules deemed difficult to synthesize by SA score may actually be straightforward to synthesize.
 728 Retrosynthesis models have much more flexibility as the building block stock and reactions can be
 729 changed, whereas SA score was designed based on the fixed PubChem corpus⁴. One could even
 730 constrain the retrosynthesis model to only include building blocks and reactions that are available
 731 in-house, similar to what was done in a collaborative work involving Pfizer⁸¹. Thus, we leave a more
 732 thorough investigation regarding SA score optimization compared to various retrosynthesis models
 733 and search algorithms for future work. In this work, only the AiZynthFinder¹¹⁻¹³ retrosynthesis
 734 model was used, which leverages Monte Carlo Tree Search and ZINC building blocks.

735 H Saturn Reward Shaping

736 This section contains details on the reward shaping functions used such that the objective functions:
737 $R_{RGFN}, R_{All\ MPO}, R_{Double\ MPO} \in [0, 1]$. Fig. H7 shows the functions for QuickVina2-GPU-
738 2.1⁵⁹⁻⁶¹ and SA score⁴. QED⁶² values were taken as is, and not subjected to reward shaping.
739 AiZynthFinder¹¹⁻¹³ returns 0 for not solved and 1 for solved. Given a molecule, all oracle evaluations
740 are aggregated via a weighted product and a single scalar value is returned as the reward:

$$R(x) = \left[\prod_i p_i(x)^{w_i} \right]^{\frac{1}{\sum_i w_i}} \quad (4)$$

741 x is a SMILES⁸², i is the index of an oracle given many oracles (MPO objective), p_i is an oracle, and
742 w_i is the weight assigned to the oracle (1 for all oracles in this work).

743 I GraphGA-augmented Experience Replay

744 Saturn⁴⁰ uses experience replay to enhance sample efficiency. GraphGA⁶⁸ can be applied on the
745 replay buffer (stores the highest rewarding molecules generated so far) by treating the replay buffer
746 as the parent population. Crossover and mutation operations then generate new molecules. For all the
747 results in this work, activating the GA decreases the AiZynthFinder solve rate relative to no GA. This
748 is because the generated molecules are not being sampled from the model itself (which is *learning*
749 to generate AiZynthFinder solvable molecules). What is gained in return is diversity recovering (as
750 found in the original Saturn⁴⁰ work). This can be advantageous since the RGFN²⁹ work defines
751 **Discovered Modes** as the number of Modes (<-10 docking score) which also have < 0.5 Tanimoto
752 similarity to every other mode. By activating the GA, more Modes are generally found, relative to no
753 GA.

754 J Saturn Batch Generation

755 Saturn⁴⁰ generates SMILES⁸² in batches of, at maximum, 16. Internally, there is an oracle caching
756 mechanism such that repeat generated SMILES are not sent for oracle evaluation, and instead, the
757 reward is retrieved from the cache. Saturn’s sample efficiency comes from the local exploration of
758 chemical space, such that, at adjacent epochs, identical SMILES can be generated. The effect is that
759 at each generation epoch, sometimes only a few *new* (not generated before) SMILES are generated.
760 In Fig. 4c, some batches have 0% solve rate by AiZynthFinder. These are batches that only have a few
761 new SMILES that happen not to be solvable. If one new SMILES is generated, it being unsolvable
762 equates to 0% solve rate.