

# Molecular Design in Synthetically Accessible Chemical Space via Deep Reinforcement Learning

Julien Horwood\* and Emmanuel Noutahi\*



Cite This: *ACS Omega* 2020, 5, 32984–32994



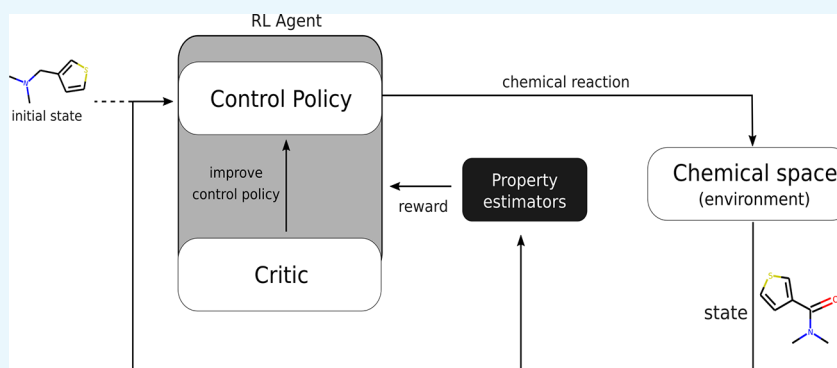
Read Online

ACCESS |

Metrics & More

Article Recommendations

Supporting Information



**ABSTRACT:** The fundamental goal of generative drug design is to propose optimized molecules that meet predefined activity, selectivity, and pharmacokinetic criteria. Despite recent progress, we argue that existing generative methods are limited in their ability to favorably shift the distributions of molecular properties during optimization. We instead propose a novel Reinforcement Learning framework for molecular design in which an agent learns to directly optimize through a space of synthetically accessible drug-like molecules. This becomes possible by defining transitions in our Markov decision process as chemical reactions and allows us to leverage synthetic routes as an inductive bias. We validate our method by demonstrating that it outperforms existing state-of-the-art approaches in the optimization of pharmacologically relevant objectives, while results on multi-objective optimization tasks suggest increased scalability to realistic pharmaceutical design problems.

## 1. INTRODUCTION

Following advances in generative modeling for domains such as computer vision and natural language processing, there has been increased interest in applying generative methods to drug discovery. However, such approaches often fail to address numerous technical challenges inherent to molecular design, including accurate molecular reconstruction, efficient exploration of chemical space, and synthetic tractability of generated molecules. Further, these approaches bias the generation of molecules toward the data distribution over which they were trained, restricting their ability to discover truly novel compounds. Previous work<sup>1,2</sup> has attempted to address these issues by framing molecular design as a Reinforcement Learning (RL) problem<sup>3</sup> in which an agent learns a mapping from a given molecular state to atoms that can be added to the molecule in a stepwise manner. These approaches generally ensure the validity of the generated compounds and avoid the need to learn a latent space mapping from the data. However, they do not address the issue of synthetic tractability, and the proposed atom-by-atom environment transitions prevent rapid exploration of chemical space.

We instead approach the problem in a way that incorporates a favorable bias into the Markov decision process. Specifically, we

define the environment's state transitions as sequences of chemical reactions, allowing us to address the common issue of synthetic accessibility. While ensuring synthesizability of computationally generated ligands is challenging, our framework treats synthesizability as a feature rather than as a constraint. Our approach, deemed REACTOR (Reaction-driven Objective Reinforcement), thus addresses a common limitation of existing methods whereby the synthetic routes for generated molecules are unknown and require challenging retrosynthetic planning. Importantly, the REACTOR framework is able to efficiently explore synthetically accessible chemical space in a goal-directed manner, while also providing a theoretically valid synthetic route for each generated compound.

We benchmark our approach against previous methods, focusing on the task of identifying novel ligands for the D2

Received: August 27, 2020

Accepted: October 27, 2020

Published: December 15, 2020



dopamine receptor, a G protein-coupled receptor involved in a wide range of neuropsychiatric and neurodegenerative disorders.<sup>4</sup> In doing so, we find that our approach outperforms previous state-of-the-art methods, is robust to the addition of multiple optimization criteria, and produces synthetically accessible, drug-like molecules by design.

## 2. RELATED WORK

Computational drug design has traditionally relied on domain knowledge and heuristic algorithms. Recently, however, several machine learning-based generative approaches have also been proposed. Many of these methods, such as ORGAN,<sup>5</sup> take advantage of the SMILES representation using recurrent neural networks (RNNs) but have difficulties generating syntactically valid SMILES. Graph-based approaches<sup>6–8</sup> have also been proposed and generally result in improved chemical validity. These methods learn a mapping from molecular graphs to a high-dimensional latent space from which molecules can be sampled and optimized. In contrast, pure Reinforcement Learning algorithms such as<sup>1,2</sup> treat molecular generation as a Markov decision process in which molecules are assembled from basic building blocks such as atoms and fragments. However, a core limitation of existing methods is the forward-synthetic feasibility of proposed designs. To overcome these limitations, Button et al.<sup>9</sup> proposed a hybrid rule-based and machine learning approach in which molecules are assembled from fragments under synthetic accessibility constraints in an iterative single-step process. However, this approach is limited in terms of the flexibility of its optimization objectives as it only allows for generation of molecules similar to a given template ligand.

In order to have practical value, methods for computational drug design must also make appropriate tradeoffs between molecular generation, which focuses on the construction of novel and valid molecules, and molecular optimization, which focuses on the properties of the generated compounds. While prior work has attempted to address these challenges simultaneously, this can lead to sub-optimal results by favoring either the generation or optimization tasks. Generative models generally do not scale well to complex property optimization problems as they attempt to bias the generation process toward a given objective within the latent space while simultaneously optimizing over the reconstruction loss. These objectives are often conflicting, making goal-directed optimization difficult and hard to scale when multiple reward signals are required. This is generally the case in drug design where drug candidates must show activity against a given target as well as favorable selectivity, toxicity, and pharmacokinetic properties.

In contrast, atom-based Reinforcement Learning addresses the generative problem via combinatorial enumeration of molecular states<sup>2</sup> or *a posteriori* verification of molecules.<sup>1</sup> These solutions are often slow and create a bottleneck in the environment's state transitions that limits effective optimization.

## 3. METHODOLOGY

In this work, we decompose generation and optimization by delegating each problem to a distinct component of our computational framework. Specifically, we allow an environment module to handle the generative process using known chemistry as a starting point for its design, while an agent learns to effectively optimize compounds through interactions with this environment. By disambiguating the responsibilities of each component and formalizing the problem as Markov decision

processes (MDPs), we allow the modules to work *symbiotically*, exploring the chemical space both more efficiently and more effectively.

We begin with a short overview of Markov decision processes and actor-critic methods for Reinforcement Learning before defining our framework in detail.

**3.1. Background.** **3.1.1. Markov Decision Processes.** A Markov decision process (MDP)<sup>10</sup> is a powerful computational framework for sequential decision-making problems. An MDP is defined via the tuple  $(\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{P})$  where  $\mathcal{S}$  defines the possible states,  $\mathcal{A}$  denotes the possible actions that may be taken at any given time,  $\mathcal{R}$  denotes the reward distribution of the environment, and  $\mathcal{P}$  defines the dynamics of the environment. Interactions within this framework give rise to trajectories of the form  $(s_0, a_0, r_1, s_1, a_1, \dots, r_T, s_T)$  with  $T$  being the terminal time step. Crucially, an MDP assumes that

$$p(s_{t+1}, r_{t+1} | a_t, s_t, r_t, \dots, a_0, s_0) = p(s_{t+1}, r_{t+1} | a_t, s_t) \quad (1)$$

where  $t$  denotes discrete time steps.

This definition states that all prior history of a decision trajectory can be encapsulated within the preceding state, allowing an agent operating within an MDP to make decisions based solely on the current state of the environment. This assumption provides the basis for efficient learning and holds under our proposed framework. An agent's mapping from any given state to action probabilities is termed a policy, and the probability of an action  $a \in \mathcal{A}$  at state  $s$  is denoted  $\pi(a | s)$ .


**3.1.2. Policy Optimization.** The underlying objective of a Reinforcement Learning agent operating in an MDP is to optimize its policy to maximize the expected return from the environment until termination at time  $T$ , defined for any step  $t$  by

$$E_{\pi}[G_t] = E_{\pi} \left[ \sum_{m=t+1}^T \gamma^{m-t-1} r_m \right] \quad (2)$$

where  $\gamma$  is a discount factor determining the value of future rewards and the expectation is taken over the experience induced by the policy's distribution. Several approaches exist for learning a policy that maximizes this quantity. In value-based approaches,  $Q$  values of the form  $Q: \mathcal{S} \times \mathcal{A} \rightarrow \mathbf{R}$  are trained to estimate the scalar value of action–value pairs as estimates of the expected return. A policy is then derived from these values through strategies such as  $\epsilon$ -greedy control.<sup>3</sup> Alternatively, policy-based approaches attempt to parameterize the agent's behavior directly, for example, through a neural network, to produce  $\pi_{\theta}(a | s)$ . While our framework is agnostic to the specific algorithm used for learning, we choose to validate our approach with an actor-critic architecture.<sup>11</sup> This approach combines the benefits of learning a policy directly using a policy network  $\pi_{\theta}$  with a variance-reducing value network  $v_{\theta}$ . Specifically, we use a synchronous version of A3C,<sup>12</sup> which is amenable to high parallelization and further gains in training efficiency. The advantage actor-critic (A2C) objective function at time  $t$  is given by

$$L(\theta, \theta') = \log(\pi_{\theta}(a_t | s_t)) \sum_{i=0}^k (r_{t+i} + \gamma^i v_{\theta'}(s_{t+i}) - v_{\theta'}(s_t)) + \beta \mathcal{H}(\pi_{\theta}(s_t)) \quad (3)$$

$$L(\theta, \theta') = \log(\pi_{\theta}(a_t | s_t)) A_t(s_t, a_t, \theta', k) + \beta \mathcal{H}(\pi_{\theta}(s_t)) \quad (4)$$



**Figure 1.** A trajectory taken by the REACTOR agent during the optimization of affinity for the dopamine receptor D2. This trajectory provides a high-level overview of a possible synthesis route for the proposed molecule in three steps: (1) a Mitsunobu reaction, (2) a reductive amination, and (3) a Buchwald–Hartwig amination. We note that, although the proposed route is theoretically feasible, it would not be the first choice for synthesis and can easily be optimized. Nevertheless, it remains an important indication of synthesizability. We also note here that the agent learns a policy that produces structures containing a pyrrolidine/piperidine moiety, which have been shown as actives against dopamine receptors.<sup>14,15</sup>

Intuitively, maximization of eq 4's first term involves adjusting the policy parameters to align the high probability of an action with a high expected return, while the second term serves as an entropy regularizer preventing the policy from converging too quickly to sub-optimal deterministic policies.

**3.2. Molecular Design via Synthesis Trajectories.** A core insight of our framework is that we can embed knowledge about the dynamics of chemical transitions into a Reinforcement Learning system for guided exploration. In doing so, we induce a bias over the optimization task, which, given its close correspondence with natural molecular transitions, should increase learning efficiency while leading to better performance across a larger, pharmacologically relevant chemical subspace.

We propose embedding this bias into the transition model of an MDP by defining possible transitions as true chemical reactions. In doing so, we gain the additional benefit of built-in synthetic accessibility in addition to immediate access to a synthesis route for generated compounds. Lack of synthesizability is a known constraint of prior generative approaches in molecular design.<sup>13</sup> The REACTOR approach addresses this constraint by embedding synthesizability directly into the framework, leveraging synthetic routes as an inductive bias. This is demonstrated in Figure 1 where a sample trajectory is provided by REACTOR for a DRD2-optimized molecule, while a high-level overview of our framework is presented in Figure 2.

**3.2.1. Framework Definition.** We define each component of our MDP as follows.

**3.2.1.1. State Space  $\mathcal{S}$ .** We allow for any valid molecule to comprise a state in our MDP. Practically, the state space is defined as  $\{f(m) \mid m \in \mathcal{M}\}$  with  $f$  being a feature extraction function and  $\mathcal{M}$  being the space of molecules reachable given a set of chemical reactions, initialization molecules, and available reactants. We use Morgan fingerprints<sup>16</sup> with bit-length 2048 and a radius of 2 to extract feature vectors from molecules. These

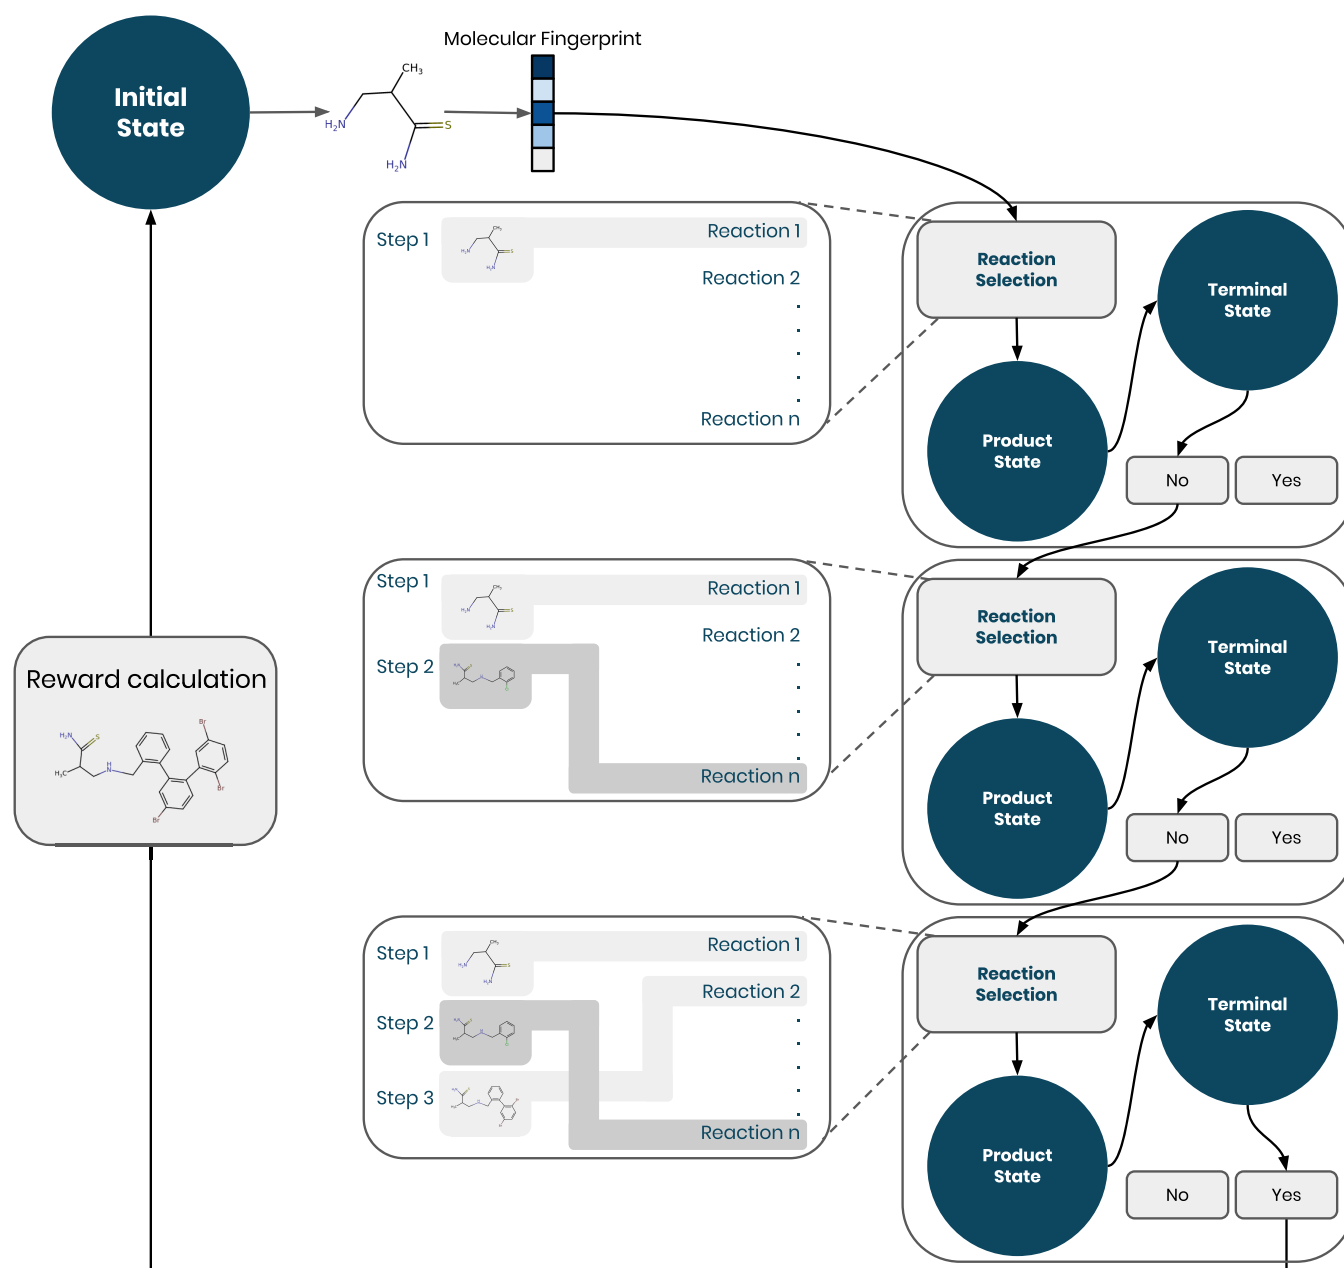
representations have been shown to provide robust and efficient featurizations, while more computationally intensive approaches like graph neural networks are yet to demonstrate significant representational benefits.<sup>17,18</sup>

**3.2.1.2. Action Space  $\mathcal{A}$ .** In its general formulation, the action space of our framework is defined hierarchically, enabling the potential application of novel approaches for hierarchical reinforcement learning. Specifically, we define a set of higher-level actions  $\mathcal{A}_o$  as a curated list of chemical reaction templates, taking the form

$$R = r_1 + r_2 + \dots + r_k \rightarrow (p_1, \dots, p_m) \quad (5)$$

Each  $r_i$  corresponds to a reactant, while each  $p_j$  is a product of this reaction. We make use of the SMARTS syntax<sup>19</sup> to represent these objects as regular expressions. We append to the high-level actions a terminal action, allowing the agent to learn to terminate an episode when the current state is deemed optimal for the objective. At step  $t$ , the state  $s_t$  thus corresponds to a single reactant in any given reaction. It is necessary to select which molecular blocks should fill in the remaining pieces for a given state and reaction selection. This gives rise to a set of primitive actions,  $\mathcal{A}_i$ , corresponding to a list of reactants derived from the reaction templates, which we also refer to as chemical building blocks. In contrast with previous methods,<sup>1,2</sup> which establish a deterministic start state such as an empty molecule or carbon atom, we initialize our environment with a randomly sampled building block that matches at minimum one reaction template. This ensures that a trajectory can take place and encourages the learned policies to be generalizable across different regions in the chemical space.

For our experiments, we work with two-reactant reaction templates and select missing reactants based on those that will most improve the next state's reward. We also select the chemical product in this manner when more than one product is



**Figure 2.** Overview of the REACTOR framework. Each episode is initialized with a molecular building block. At each step, the current state is converted to its fingerprint representation and the policy model selects a reaction to be performed. A reactant selection heuristic completes the reaction to generate the next state in the episode, while a reward of 0 is returned. Instead, if the terminal action is selected, the current state is considered as the final molecule and its reward is used to update the policy's parameters.

**Table 1. Goal-Directed Molecule Design**

objective	method	total actives	mean activity	diversity	scaff. similarity	uniqueness
DRD2	BLOCKS	3 ± 0	0.03 ± 0	0.94 ± 0	N/A <sup>a</sup>	1.0 ± 0.0
	hill climbing	43.0 ± 2.94	0.43 ± 0.03	0.878 ± 0.01	0.124 ± 0.0	1.0 ± 0.0
	ORGAN	5.333 ± 0.47	0.093 ± 0.01	0.86 ± 0.01	0.577 ± 0.11	0.873 ± 0.01
	JTVAE	4.0 ± 0.82	0.014 ± 0.0	0.934 ± 0.0	0.097 ± 0.0	0.976 ± 0.01
	GCPN	0.0 ± 0.0	0.0 ± 0.0	0.906 ± 0.0	0.12 ± 0.0	1.0 ± 0.0
	MolDQN	9.667 ± 0.47	0.816 ± 0.08	0.6 ± 0.02	N/A	0.12 ± 0.02
	REACTOR	77.0 ± 4.32	0.77 ± 0.04	0.702 ± 0.02	0.133 ± 0.01	1.0 ± 0.0

<sup>a</sup>Computation of the scaffold similarity requires the presence of a ring system, thus the N/A.

generated. Doing so collapses our hierarchical formulation into a standard MDP formulation with the reaction selection being the

only decision point. Additionally, it is likely that, for any step  $t$ , the set of possible reactions is smaller than the full action space.



In order to increase both the scalability of our framework (by allowing for larger reaction lists) and the speed of training, we use a mask over unfeasible reactions. This avoids the need for the agent to learn the chemistry of reaction feasibility and reduces the effective dimension of the action space at each step. We compare policy convergence when using a masked action space to a regular action space formulation in Figure S1. The policy then takes the form  $\pi(a_t | s_t, M(s_t, \mathbf{R}))$  with  $M$  being the environment's masking function and  $\mathbf{R}$  being the list of reaction templates.

**3.2.1.3. Reward Distribution  $\mathcal{R}$ .** Appropriate reward design is crucial given that it drives the policy optimization process. In graph convolutional policy networks,<sup>1</sup> intermediate and adversarial rewards are introduced in order to enforce drug-likeness and validity of generated compounds. In MolDQN,<sup>2</sup> these requirements are ignored, and while optimization performance increases, desirable pharmaceutical properties are often lost. In the REACTOR framework, the separation between the agent and the environment allows us to maintain property-focused rewards that guide optimization while ensuring that chemical constraints are met via environment design.

We use a deterministic reward function based on the property to be optimized. In Table 1, this corresponds to the binary prediction of compounds binding to the D2 dopamine receptor (DRD2). In Table S1, these are the penalized calculated octanol–water partition coefficient ( $c\text{Log}P$ ) and quantitative estimate of drug-likeness (QED).<sup>20</sup> In order to avoid artificially biasing our agent toward greedy policies, we remove intermediate rewards and provide evaluative feedback only at the termination of an episode. While we feel that this is a more principled view on the design process, Zhou et al.<sup>2</sup> have also suggested that, using an intermediate reward discounted by a decreasing function of the step  $t$  may improve the learning efficiency. We further apply a constraint based on the atom count of a molecule to be consistent with prior work. When molecules exceed the maximum number of atoms (38), the agent observes a reward of zero.

**3.2.1.4. Transition Model  $\mathcal{P}$ .** In the template-based REACTOR framework, state transitions are deterministic. We therefore have  $p(s_{t+1} | s_t, a_t) = 1$  according to our choice of reaction and the subsequent reactant–product selection. When modifying the reactant–selection policy, either via a stochastic heuristic such as an epsilon-greedy reactant selection or learned hierarchical policies, state transitions over the higher-level actions,  $A_o$ , become stochastic according to the internal policy's dynamics.

**3.2.2. Building Block Fragmentation.** In order to maximize the exploration capacity of the REACTOR agent, it is desirable to scale up the size of both the reaction template and reactant lists. However, current Reinforcement Learning methodology is poorly suited for very large discrete action spaces. In particular, there are approximately 76,000 building blocks available for our experiments with a wide range of possibilities matching a given reaction template position. While certain approaches propose learning a mapping from continuous to discrete action spaces,<sup>21,22</sup> we instead mitigate the dimensionality of the reactant space directly. Indeed, we leverage the BRICS<sup>23</sup> retrosynthesis rules to reduce our original reactant set to one of approximately 5000 smaller blocks. This reduces the reactant space dimension by an order of magnitude while rendering the transitions in space less extreme and thus more flexible. Additionally, we may limit the size of the set of reactants under consideration at any given step, treating this as a

hyperparameter. For our experiments, we set this to 100 reactants, finding little variation when selecting reactants in a greedy manner.

## 4. RESULTS AND DISCUSSION

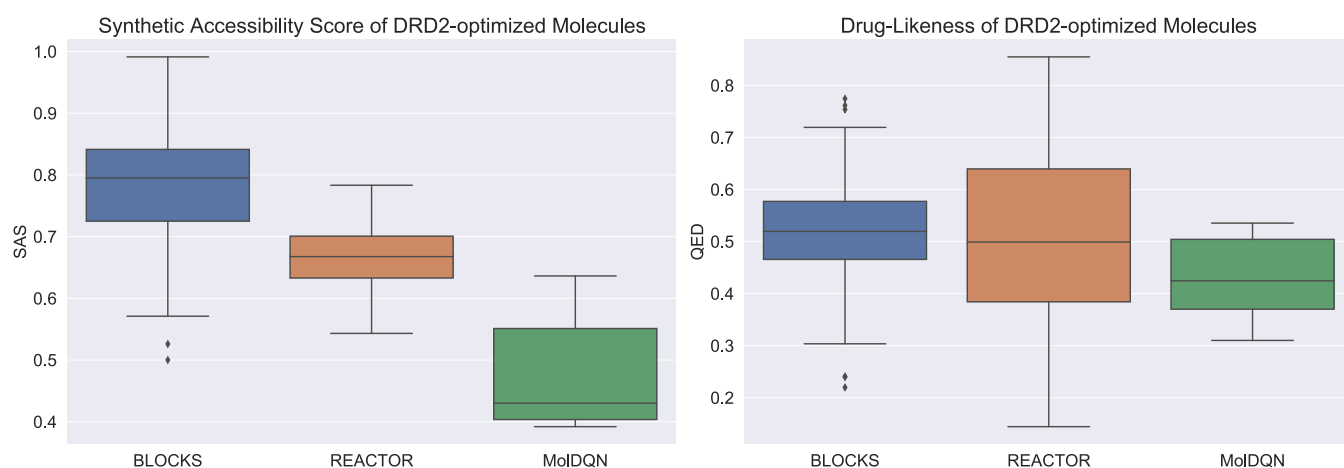
To validate our framework, we benchmark its performance on goal-directed design tasks, focusing primarily on predicted activity for the D2 dopamine receptor. We frame this objective as a sparse reward, using a binary activity indication to simulate a hit discovery setting. In order to maintain consistency with experiments done in prior work, we perform additional experiments on penalized  $c\text{Log}P$  and QED with the results presented in the Supporting Information.

In order to better understand the exploration behavior of our approach, we also investigate the nature of the trajectories generated by the REACTOR policies, showing that policies retain drug-likeness across all optimization objectives, while also exploring distinct regions of the chemical space.

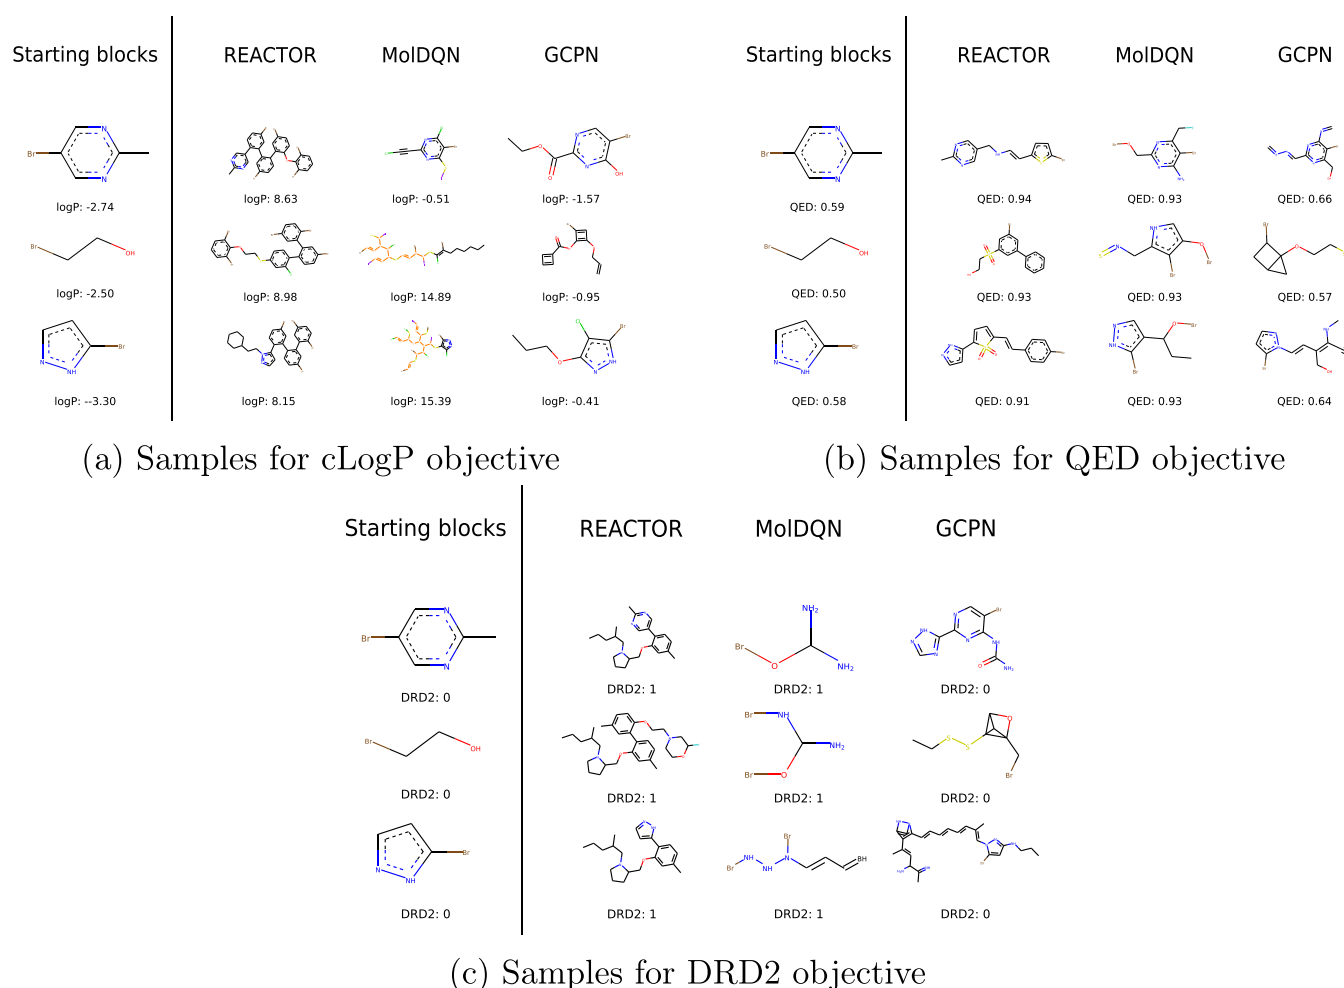
**4.1. Experimental Setup. 4.1.1. Reaction Data.** For these experiments, the set of reactions used was obtained from Konze et al.<sup>24</sup> with the final list consisting of 127 reactions following curation for specificity and validity. The set of reactants are drawn from PubChem<sup>8,25</sup> totaling 76,208 building blocks matching the reaction templates. Following the retrosynthesis methodology introduced above, these lists were reduced to approximately 5000 smaller reactants with 90 reaction templates matching these blocks. This allows us to make the space of action possibilities more tractable while rendering the exploration of chemical space more flexible due to each transition corresponding to smaller steps in space. Naturally, this action space does not encompass all chemical transformations, which may be of interest in a general setting. However, it is straightforward to extend the reaction templates and associated building blocks to tailor the search space to the data available for a given design objective.

**4.1.2. Empirical Reward Models.** While generative models are biased by their data distribution, RL-driven molecule design may be biased implicitly by training data used for an empirical reward model. Thus, it is crucial that these models provide robust generalization. A model that is overly simplistic, as is seen for the  $c\text{Log}P$  experiments, may lead to agents exploiting particular biases, leading to pharmacologically undesirable molecules. Training details for the DRD1, DRD2, DRD3, and Caco-2 models are found in the Supporting Information.

**4.1.3. Baselines.** We compare our approach to two recent methods in deep generative molecular modeling, JT-VAE and ORGAN.<sup>5,8</sup> Each of these approaches was first pre-trained for up to 48 h on the same compute facility, a single machine with 1 NVIDIA Tesla K80 GPU and 16 CPU cores. Property optimization was then performed using the same procedures as described in the original papers. We also compare our method with two state-of-the-art Reinforcement Learning approaches, graph-convolutional policy networks and MolDQN.<sup>1,2</sup> Each algorithm was run using the open-sourced code from the authors, while we enforced the same reward function implementation across methods to ensure consistency. We ran the GCPN using 32 CPU cores for approximately 24 h (against 8 h in the original paper) and MolDQN for 20,000 episodes (against 5000 episodes in the original paper). In addition, we added a steepest-ascent hill-climbing baseline using the REACTOR environment to demonstrate that for simple mostly greedy objectives such as  $c\text{Log}P$  and QED, simple search policies may provide reasonable performance. In contrast, learned



**Figure 3.** Synthetic accessibility and drug-likeness score distributions of molecules optimized for DRD2 and the starting blocks.



**Figure 4.** (a–c) Sample molecules produced for each objective by each RL algorithm.

traversals of space become necessary for complex tasks such as DRD2.

**4.1.4. Evaluation.** Given the inherent differences between generative and reinforcement learning models, evaluation was adapted to remain consistent within each class of algorithms. JT-VAE and ORGAN were evaluated based on decoded samples from their latent space using the best results across training checkpoints with statistics for JT-VAE computed over three random seeds. Given the prohibitive cost of training ORGAN,

results are given over a single seed and averaged over three sets of 100 samples. Other baselines were compared based on three sets of 100 building blocks used as starting states. Statistics are reported over sets, while the statistics of the initial states are shown by BLOCKS.

We prioritize the evaluation of each method based on the total number of active molecules identified, as determined by the environment reward model, given that this corresponds most to the underlying objective of *de novo* design. Indeed, in a hit

discovery scenario, a user may be most interested in identifying the maximal number of unique potential hits, leaving potency optimization to later stages in the lead optimization process. We denote this quantity by “total actives” in Table 1. “mean activity” corresponds to the percentage of generated molecules that are predicted to be active for the DRD2 receptor. In both Table 1 and Table S1, the mean reward (“mean activity”) was computed based on the set of unique molecules generated by each algorithm in order to avoid artificially favoring methods, which often generate the same molecule. Diversity corresponds to the average pairwise Tanimoto distance among generated molecules, while “scaff. similarity” corresponds to the average pairwise similarity between the scaffolds of the compounds, as implemented by the MOSES repository.<sup>26</sup> Finally, we limited the number of atoms to 38 for all single-objective tasks, as done in prior work,<sup>1,2,8</sup> and 50 for the multi-objective tasks.

**4.2. Goal-Directed *de Novo* Design.** Results on the unconstrained design task show that REACTOR identifies the most active molecules for the DRD2 objective. Furthermore, we observe that REACTOR maintains high diversity and uniqueness in addition to robust performance. This is a crucial characteristic as it implies that the agent is able to optimize the space surrounding each starting molecule without reverting to the same molecule to optimize the scalar reward signal. In Table S1, REACTOR also achieves a higher reward on QED, while remaining competitive on penalized *cLogP* despite the simplistic nature of this objective favoring atom-by-atom transitions. We note that, while MolDQN exhibits higher mean activity, this is attributed to the fact that the optimization tends to collapse into generating the same molecule. This explains why the total number of active molecules identified remains low despite the mean activity suggesting good performance on the task.

Training efficiency is an important practical consideration while deploying methods for *de novo* design. Generative models first require learning a mapping of molecules to the latent space before training for property optimization. During our experiments, this resulted in more than 48 h of training time, and after which, training was stopped. Reinforcement Learning methods trained faster but generally failed to converge within 24 h. We ran MolDQN for 20,000 episodes, taking between 24 and 48 h, while GCPN was stopped after 24 h on 32 CPU cores. In contrast, our approach converges within approximately 2 h of training on 40 CPU cores for the *cLogP* and QED objectives while consuming less memory than GCPN for 32 cores and terminates under 24 h for the D2-related tasks. In order to make effective use of parallelization, we leveraged the implementation of A2C provided by the rllib library.<sup>27</sup>

**4.3. Synthetic Tractability and Desirability of Optimized Compounds.** Given the narrow perspective offered by quantitative benchmarks for molecular design models,<sup>26</sup> it is equally important to qualitatively assess the behavior of these models by examining generated compounds. Figure 4 provides samples generated by each RL method across all objectives. Since the computational estimation of *cLogP* relies on the Wildman–Crippen method,<sup>28</sup> which assigns a high atomic contribution to halogens and phosphorous, the atom-based action space of MolDQN produces samples that are heavily biased toward these atoms, resulting in molecules that are well optimized for the task but neither synthetically-accessible nor drug-like. This generation bias was not observable in previously reported benchmarks where atom types were only limited to carbon, oxygen, nitrogen, sulfur, and halogens.<sup>2</sup> Furthermore,

MolDQN samples for the DRD2 task lack a ring system, and whereas molecules from GCPN have one, none adequately optimizes for the objective.

In contrast, REACTOR appears to produce more pharmacologically desirable compounds without explicitly considering this as an optimization objective. This is supported by Figure 3, which illustrates the shift in synthetic accessibility scores<sup>29</sup> and drug-likeness for the DRD2-active molecules produced by REACTOR and MolDQN. This suggests that REACTOR is able to simultaneously solve the DRD2 task while maintaining favorable distributions for synthetic accessibility and drug-likeness.

Further, as shown in Figures 1 and 7, optimized compounds are provided along with a possible route of synthesis. While such trajectories may not be optimal given that they are limited by the reward design and the set of reaction templates available, they provide a crucial indication of synthesizability. Further, it is possible to encourage trajectories to be more efficient by limiting the number of synthesis steps per episode or by incorporating additional costs such as reactant availability and synthesis difficulty in the reward design. In certain applications, it may also be desirable to increase specificity of the reaction templates via group protection. Gao and Coley<sup>13</sup> detail the lack of consideration for synthetic tractability in current molecular optimization approaches, highlighting that this is a necessary requirement for application of these methods in drug discovery. While alternative ideas aiming to embed synthesizability constraints into generative models have recently been explored,<sup>9,30,31</sup> REACTOR is the first approach that explicitly addresses synthetic feasibility by optimizing directly in the space of synthesizable compounds using Reinforcement Learning.

**4.4. Multi-objective Optimization.** Practical methods for computational drug design must be robust to the optimization of multiple properties. Indeed, beyond the agonistic or antagonistic effects of a small molecule, properties such as the selectivity, solubility, drug-likeness, and permeability of a drug candidate must be considered. To validate the REACTOR framework under this setting, we consider the task of optimizing for selective DRD2 ligands. Dopamine receptors are grouped into two classes: D1-like receptors (DRD1 and DRD5) and D2-like receptors (DRD2, DRD3, and DRD4). Although these receptors are well studied, design of drugs selective across subtypes remains a considerable challenge. In particular, as DRD1 and DRD3 share 78% structural similarity in their transmembrane region,<sup>4,32</sup> it is very challenging to identify small molecules that can selectively bind to them and modulate their activity. We therefore assess the performance both on selectivity across classes (using DRD1 as an off-target) and within classes (using DRD3 as an off-target). We then analyze how our framework performs as we increase the number of design objectives. For these experiments, we focus our comparison on MolDQN as it outperforms other state-of-the-art methods on the single-objective tasks. Our approach in combining multiple objectives is that of reward scalarization. Formally, a vector of reward signals is aggregated via a mapping,  $\mathcal{S}: R_1 \times \dots \times R_k \rightarrow \mathbf{R}$ , thus collapsing the multi-objective MDP<sup>33</sup> into a standard MDP formulation. While the simplest and most common approach to scalarization is to use a weighted sum of the individual reward signals, we adopt a Chebyshev scalarization scheme<sup>34</sup> whereby reward signals are aggregated via the weighted Chebyshev metric:

Table 2. DRD2 Selectivity

objective	method	total actives	mean reward	diversity	scaff. similarity	uniqueness
D2/D1	MolDQN	9.0 ± 1.41	0.64 ± 0.07	0.502 ± 0.01	N/A	0.14 ± 0.01
	REACTOR	36.667 ± 4.99	0.368 ± 0.05	0.599 ± 0.01	0.139 ± 0.01	0.997 ± 0.0
D2/D3	MolDQN	25.667 ± 3.09	0.884 ± 0.07	0.746 ± 0.05	N/A	0.29 ± 0.02
	REACTOR	53.0 ± 8.29	0.53 ± 0.08	0.692 ± 0.03	0.147 ± 0.01	1.0 ± 0.0

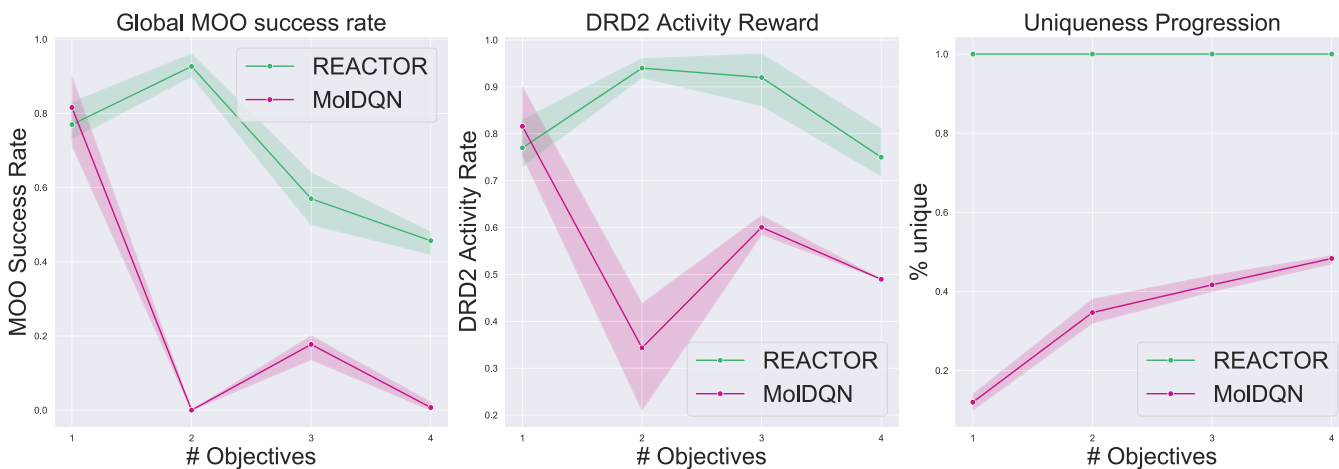


Figure 5. Reward progression as the number of optimization objectives increases.

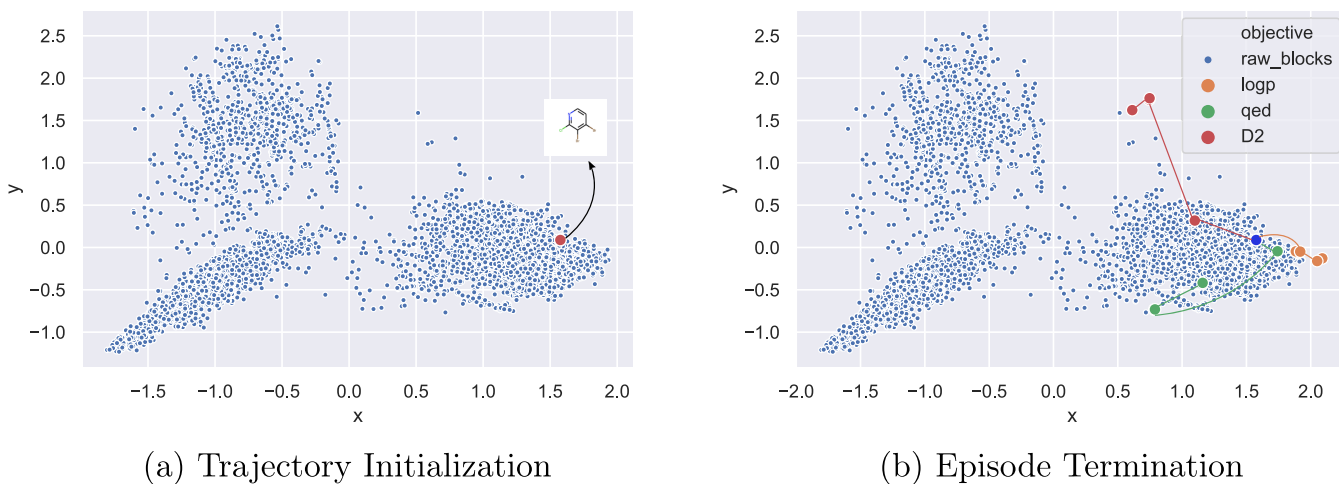


Figure 6. (a, b) Trajectory steps of the REACTOR agent for each objective, starting with the same building block. The RL agent shifts the molecule toward different regions in space to identify the relevant local maximum.

$$r = -\max_i w_i (|r_i - z_i^*|) \quad (6)$$

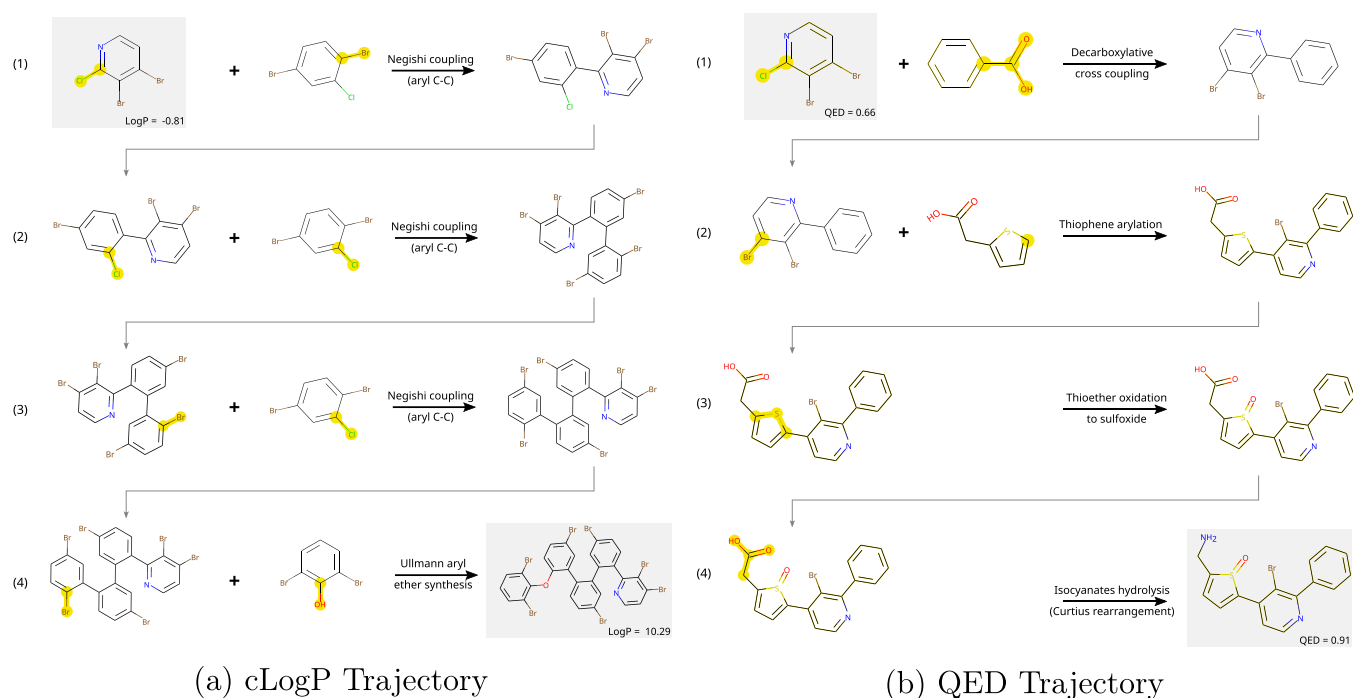
where  $\vec{z}^*$  is a utopian vector,  $\vec{w}$  is assigned to the relative preferences for each objective, and  $i$  indexes the objectives. For our experiments, we consider rewards that are constrained to a range between 0 and 1 such that the utopian point is always  $\vec{1}$ , rendering the dynamics of each reward signal more comparable, and we assign equal preferences to the objectives. For the selectivity tasks, given that both rewards are binary, we use a soft version of this scalarization scheme corresponding the negative Euclidean distance to the optimal point. This allows the reward signal to differentiate between reaching 0,1 and both of the objectives. While Chebyshev scalarization was introduced for the setting of tabular Reinforcement Learning, we may interpret it in the function approximation setting as defining an adaptive

curriculum whereby the optimization focus shifts dynamically according to the objective most distant from  $\vec{z}^*$ .

**4.4.1. DRD2 Selectivity.** The total number of actives in Table 2 corresponds to the number of unique molecules that were found to satisfy all objectives, while the mean reward in Table 2 and Figure 5 is computed as the proportion of evaluation episodes for which the algorithms optimize all desired objectives. In Table 2, we find that REACTOR maintains the ability to identify a higher number of desirable molecules on the selectivity tasks, optimizing for DRD2 while avoiding off-target activity on the D1 and D3 receptors. Further, it is able to outperform MolDQN while maintaining very low scaffold similarity among generated molecules.

**4.4.2. Robustness to Multiple Objectives.** In addition to off-target selectivity, we assess the robustness of each method's performance as we increase the number of pharmacologically





**Figure 7.** (a, b) Trajectories taken by the REACTOR agent from the same building block for different objectives. Note that the reaction steps are simplified and are mainly indicative of synthesizability. For example, the Negishi coupling reaction would first require the formation of an organozinc precursor. Furthermore, selectivity is low at some steps, which will result in a mixture of products, unless reacting groups are protected.

relevant property objectives to optimize. Specifically, we compare the following combinations of rewards:

- DRD2 with range-targeted *cLogP* (2 objectives) according to the Ghose filter<sup>35</sup>
- DRD2, range *cLogP*, and a molecular weight ranging from 180 to 600 (3 objectives)
- DRD2, range *cLogP*, target molecular weight, and drug absorption as indicated by a model trained on data for the Caco-2 permeability assay<sup>36</sup> (4 objectives)

For the range-targeted *cLogP*, molecular weight, and permeability objectives, the component-wise reward is 0 when the molecule falls within the desired range. Otherwise, the distance to the objective is mapped to a range of (0,1]. Given that the DRD2 objective is binary, this implicitly prioritizes the optimization for this reward.

Figure 5 shows that REACTOR demonstrates greater robustness to an increasing number of design objectives. Indeed, while both methods see diminishing success rates in optimizing for multiple objectives, the performance of REACTOR diminishes gradually, while MolDQN's performance collapses. Furthermore, REACTOR maintains the ability to generate unique terminal states throughout.

**4.5. Goal-Directed Exploration.** In order to gain further insight into the nature of the trajectories generated by the REACTOR agent, we plotted two alternative views of optimization routes generated for the same building block across each single-property objective. In Figure 6, we fit a principal component analysis (PCA)<sup>37</sup> on the space of building blocks to identify the location of the initial state and subsequently transform the next states generated by the RL agent onto this space. We find that the initial molecule is clearly shifted to distinct regions in space, while the magnitude of the transitions suggest efficient traversal of the space. This provides further evidence that exploration through space is a function of

reward design and is mostly unbiased by the data distribution of initialization states. Figure 7 shows the same trajectories with their corresponding reactions and intermediate molecular states. We find that optimized molecules generally contain the starting structure. We believe this to be a desirable property given that real-life design cycles are often focused on a fixed scaffold or set of core structures. We also note that the policy learned by our REACTOR framework is able to generalize over different starting blocks, suggesting that it achieves generation of structurally diverse and novel compounds.

## 5. CONCLUSIONS

This work proposes a novel approach to molecular design that defines state transitions as chemical reactions within a Reinforcement Learning framework. We demonstrate that our framework leads to globally improved performance, as measured by reward and diversity of generated molecules, as well as greater training efficiency while producing more drug-like molecules. Analysis of REACTOR's robustness to multiple optimization criteria, coupled with its ability to maintain predicted activity on the DRD2 receptor, suggests increased potential for successful application in drug discovery. Furthermore, molecules generated by this framework exhibit better synthetic accessibility by design with one viable synthesis route also suggested. Although the reactivity and stability of the optimized molecules remain a potential issue, REACTOR's efficiency in a multiple optimization setting suggests that this can be addressed by explicitly considering them as additional design objectives.

Future work aims to build on this framework by making use of its hierarchical formulation to guide agent policies both at the higher reaction and lower reactant levels, exploring proposals from h-DQN<sup>38</sup> for hierarchical value functions, or the option-critic framework<sup>39</sup> as a starting point. We also plan to expand the effective state space of our MDP by embedding a synthesis model with transformer-based architectures showing promise<sup>40</sup>

as the MDP transition model. Because practical *de novo* design requires optimization of multiple criteria simultaneously, we believe that the efficiency of our design framework provides a robust foundation for such tasks and hope to expand on existing approaches<sup>41–43</sup> for multi-objective Reinforcement Learning. Finally, we intend to validate the proposed synthetic routes and bio-activity of generated molecules experimentally to better demonstrate real-world utility.

## ■ ASSOCIATED CONTENT

### SI Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acsomega.0c04153>.

Results on QED/cLogP optimization tasks (Table S1); ablation study for use of a masked action space (Figure S1); molecule samples from REACTOR: DRD2 task, DRD2 activity/DRD1 selectivity task, and DRD2 activity/DRD3 selectivity task (Figure S2); and model training details: DRD2 model report (Figure S3) and DRD1/DRD3/Caco-2 model training details (PDF)

## ■ AUTHOR INFORMATION

### Corresponding Authors

**Julien Horwood** – *InVivo AI, Montreal, Quebec H2S 3H1, Canada; Mila, Université de Montréal, Montreal, Quebec H2S 3H1, Canada; [orcid.org/0000-0001-8987-7164](https://orcid.org/0000-0001-8987-7164); Email: [julien@invivoai.com](mailto:julien@invivoai.com)*

**Emmanuel Noutahi** – *InVivo AI, Montreal, Quebec H2S 3H1, Canada; Email: [emmanuel@invivoai.com](mailto:emmanuel@invivoai.com)*

Complete contact information is available at:

<https://pubs.acs.org/doi/10.1021/acsomega.0c04153>

### Notes

The authors declare the following competing financial interest(s): The authors are employed by InVivo AI, which is actively involved in the pursuit of novel methodologies for drug discovery.

## ■ ACKNOWLEDGMENTS

The authors thank Daniel Cohen, Sébastien Giguère, Violet Guo, Michael Craig, Connor Coley, and Grant Wishart for reviewing the manuscript and for their helpful comments. The methods and algorithms presented here were developed at InVivo AI.

## ■ ADDITIONAL NOTE

<sup>§</sup>A mapping of SMILES to PubChem ID is available upon request.

## ■ REFERENCES

- (1) You, J.; Liu, B.; Ying, Z.; Pande, V.; Leskovec, J. Graph Convolutional Policy Network for Goal-Directed Molecular Graph Generation. In *Advances in Neural Information Processing Systems*; Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R., Eds.; Curran Associates, Inc.: 2018, pp. 6410–6421.
- (2) Zhou, Z.; Kearnes, S.; Li, L.; Zare, R. N.; Riley, P. Optimization of Molecules via Deep Reinforcement Learning. *Sci. Rep.* **2019**, *9*, 10752.
- (3) Sutton, R. S.; Barto, A. G. *Reinforcement learning: An introduction*; MIT press: 2018.
- (4) Moritz, A. E.; Free, R. B.; Sibley, D. R. Advances and challenges in the search for D<sub>2</sub> and D<sub>3</sub> dopamine receptor-selective compounds. *Cell. Signalling* **2018**, *41*, 75–81.

- (5) Guimaraes, G. L.; Sanchez-Lengeling, B.; Outeiral, C.; Farias, P. L. C.; Aspuru-Guzik, A. Objective-Reinforced Generative Adversarial Networks (ORGAN) for Sequence Generation Models. *arXiv:1705.10843 [cs, stat]* **2018**, arXiv: 1705.10843.

- (6) De Cao, N.; Kipf, T. MolGAN: An implicit generative model for small molecular graphs. *arXiv ICML 2018 workshop on Theoretical Foundations and Applications of Deep Generative Models* **2018**.

- (7) Liu, Q.; Allamanis, M.; Brockschmidt, M.; Gaunt, A., Constrained Graph Variational Autoencoders for Molecule Design. In *Advances in Neural Information Processing Systems* 31; Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R., Eds.; Curran Associates, Inc.: 2018; pp. 7795–7804.

- (8) Jin, W.; Barzilay, R.; Jaakkola, T. Junction Tree Variational Autoencoder for Molecular Graph Generation. *arXiv:1802.04364 [cs, stat]* **2019**, arXiv: 1802.04364.

- (9) Button, A.; Merk, D.; Hiss, J. A.; Schneider, G. Automated *de novo* molecular design by hybrid machine intelligence and rule-driven chemical synthesis. *Nat. Mach. Intell.* **2019**, *1*, 307–315.

- (10) Bellman, R. A Markovian Decision Process. *J. Math. Mech.* **1957**, *6*, 679–684.

- (11) Konda, V. R.; Tsitsiklis, J. N., Actor-Critic Algorithms. In *Advances in Neural Information Processing Systems* 12; Solla, S. A., Leen, T. K., Müller, K., Eds.; MIT Press: 2000; pp. 1008–1014.

- (12) Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T.; Harley, T.; Silver, D.; Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In *Proceedings of The 33rd International Conference on Machine Learning*; ed. by Balcan, M. F.; Weinberger, K. Q., PMLR: New York, New York, USA, 2016; Vol. 48, pp. 1928–1937.

- (13) Gao, W.; Coley, C. W. The Synthesizability of Molecules Proposed by Generative Models. *J. Chem. Inf. Model.* **2020**, DOI: [10.1021/acs.jcim.0c00174](https://doi.org/10.1021/acs.jcim.0c00174).

- (14) Martelle, J. L.; Nader, M. A. A Review of the Discovery, Pharmacological Characterization, and Behavioral Effects of the Dopamine D<sub>2</sub>-Like Receptor Antagonist Eticlopride. *CNS Neurosci. Ther.* **2008**, *14*, 248–262.

- (15) Gilligan, P. J.; Cain, G. A.; Christos, T. E.; Cook, L.; Drummond, S.; Johnson, A. L.; Kergaye, A. A.; McElroy, J. F.; Rohrbach, K. W. Novel piperidine. sigma. receptor ligands as potential antipsychotic drugs. *J. Med. Chem.* **1992**, *35*, 4344–4361.

- (16) Rogers, D.; Hahn, M. Extended-Connectivity Fingerprints. *J. Chem. Inf. Model.* **2010**, *50*, 742–754.

- (17) Yang, K.; Swanson, K.; Jin, W.; Coley, C.; Eiden, P.; Gao, H.; Guzman-Perez, A.; Hopper, T.; Kelley, B.; Mathea, M.; et al. Correction to Analyzing Learned Molecular Representations for Property Prediction. *J. Chem. Inf. Model.* **2019**, *59*, 5304–5305.

- (18) Kearnes, S.; McCloskey, K.; Berndl, M.; Pande, V.; Riley, P. Molecular Graph Convolutions: Moving Beyond Fingerprints. *J. Comput.-Aided Mol. Des.* **2016**, *30*, 595–608.

- (19) Daylight Theory: SMARTS - A Language for Describing Molecular Patterns; <https://www.daylight.com/dayhtml/doc/theory/theory.smarts.html>, (accessed Jan, 2020).

- (20) Bickerton, G. R.; Paolini, G. V.; Besnard, J.; Muresan, S.; Hopkins, A. L. Quantifying the chemical beauty of drugs. *Nat. Chem.* **2012**, *4*, 90–98.

- (21) Chandak, Y.; Theocharous, G.; Kostas, J.; Jordan, S.; Thomas, P. S. Learning Action Representations for Reinforcement Learning. *arXiv:1902.00183 [cs, stat]* **2019**, arXiv: 1902.00183.

- (22) Dulac-Arnold, G.; Evans, R.; van Hasselt, H.; Sunehag, P.; Lillicrap, T.; Hunt, J.; Mann, T.; Weber, T.; Degris, T.; Coppin, B. Deep Reinforcement Learning in Large Discrete Action Spaces. *arXiv:1512.07679 [cs, stat]* **2016**, arXiv: 1512.07679.

- (23) Degen, J.; Wegscheid-Gerlach, C.; Zaliani, A.; Rarey, M. On the Art of Compiling and Using 'Drug-Like' Chemical Fragment Spaces. *ChemMedChem* **2008**, *3*, 1503–1507.

- (24) Konze, K. D.; Bos, P. H.; Dahlgren, M. K.; Leswing, K.; Tubert-Brohman, I.; Bortolato, A.; Robbason, B.; Abel, R.; Bhat, S. Reaction-Based Enumeration, Active Learning, and Free Energy Calculations To Rapidly Explore Synthetically Tractable Chemical Space and Optimize

Potency of Cyclin-Dependent Kinase 2 Inhibitors. *J. Chem. Inf. Model.* **2019**, *59*, 3782–3793.

(25) Kim, S.; Chen, J.; Cheng, T.; Gindulyte, A.; He, J.; He, S.; Li, Q.; Shoemaker, B. A.; Thiessen, P. A.; Yu, B.; Zaslavsky, L.; Zhang, J.; Bolton, E. E. PubChem 2019 update: improved access to chemical data. *Nucleic Acids Res.* **2019**, D1102.

(26) Polykovskiy, D.; Zhebrak, A.; Sanchez-Lengeling, B.; Golovanov, S.; Tatanov, O.; Belyaev, S.; Kurbanov, R.; Artamonov, A.; Aladinskiy, V.; Veselov, M.; Kadurin, A.; Nikolenko, S.; Aspuru-Guzik, A.; Zhavoronkov, A. Molecular Sets (MOSES): A Benchmarking Platform for Molecular Generation Models. *arXiv preprint arXiv:1811.12823* **2018**.

(27) Liang, E.; Liaw, R.; Nishihara, R.; Moritz, P.; Fox, R.; Goldberg, K.; Gonzalez, J.; Jordan, M.; Stoica, I. RLlib: Abstractions for Distributed Reinforcement Learning. In *Proceedings of the 35th International Conference on Machine Learning*; ed. by Dy, J.; Krause, A., PMLR: Stockholmsmässan, Stockholm Sweden, 2018; Vol. 80, pp. 3053–3062.

(28) Wildman, S. A.; Crippen, G. M. Prediction of physicochemical parameters by atomic contributions. *J. Chem. Inf. Comput. Sci.* **1999**, *39*, 868–873.

(29) Ertl, P.; Schuffenhauer, A. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *J. Cheminf.* **2009**, *1*, 8.

(30) Korovina, K.; Xu, S.; Kandasamy, K.; Neiswanger, W.; Poczos, B.; Schneider, J.; Xing, E. ChemBO: Bayesian Optimization of Small Organic Molecules with Synthesizable Recommendations. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*; ed. by Chiappa, S.; Calandra, R., PMLR: 2020; Vol. 108, pp. 3393–3403.

(31) Bradshaw, J.; Paige, B.; Kusner, M. J.; Segler, M.; Hernández-Lobato, J. M., A Model to Search for Synthesizable Molecules. In *Advances in Neural Information Processing Systems*; Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R., Eds.; Curran Associates, Inc.: 2019, pp. 7937–7949.

(32) Sibley, D. R.; Monsma, F. J., Jr. Molecular biology of dopamine receptors. *Trends Pharmacol. Sci.* **1992**, *13*, 61–69.

(33) Wiering, M. A.; de Jong, E. D. Computing Optimal Stationary Policies for Multi-Objective Markov Decision Processes. In *2007 IEEE International Symposium on Approximate Dynamic Programming and Reinforcement Learning*; IEEE: 2007; pp. 158–165.

(34) Van Moffaert, K.; Drugan, M. M.; Nowe, A. Scalarized multi-objective reinforcement learning: Novel design techniques. In *2013 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*; IEEE: 2013; pp. 191–199.

(35) Ghose, A. K.; Viswanadhan, V. N.; Wendoloski, J. J. A Knowledge-Based Approach in Designing Combinatorial or Medicinal Chemistry Libraries for Drug Discovery. 1. A Qualitative and Quantitative Characterization of Known Drug Databases. *J. Comb. Chem.* **1999**, *1*, 55–68.

(36) Van Breemen, R. B.; Li, Y. Caco-2 cell permeability assays to measure drug absorption. *Expert Opin. Drug Metab. Toxicol.* **2005**, *1*, 175–185.

(37) Pearson, K. LIII. On lines and planes of closest fit to systems of points in space. *London, Edinburgh Dublin Philos. Mag. J. Sci.* **2010**, *2*, 559–572.

(38) Kulkarni, T. D.; Narasimhan, K.; Saeedi, A.; Tenenbaum, J. Hierarchical Deep Reinforcement Learning: Integrating Temporal Abstraction and Intrinsic Motivation. In *Advances in Neural Information Processing Systems 29*; Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., Garnett, R., Eds.; Curran Associates, Inc.: 2016; pp. 3675–3683.

(39) Bacon, P.-L.; Harb, J.; Precup, D. The option-critic architecture. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*; AAAI Press: San Francisco, California, USA, 2017; pp. 1726–1734.

(40) Schwaller, P.; Laino, T.; Gaudin, T.; Bolgar, P.; Hunter, C. A.; Bekas, C.; Lee, A. A. Molecular Transformer: A Model for Uncertainty-Calibrated Chemical Reaction Prediction. *ACS Cent. Sci.* **2019**, *5*, 1572–1583.

(41) Abels, A.; Roijers, D.; Lenaerts, T.; Nowe, A.; Steckelmacher, D. Dynamic Weights in Multi-Objective Deep Reinforcement Learning. In *Proceedings of the 36th International Conference on Machine Learning*; ed. by Chaudhuri, K.; Salakhutdinov, R., PMLR: 2019; pp. 11–20.

(42) Moffaert, K. V.; Nowe, A. Multi-Objective Reinforcement Learning using Sets of Pareto Dominating Policies. *J. Mach. Learn. Res.* **2014**, *15*, 3663–3692.

(43) Yang, R.; Sun, X.; Narasimhan, K. A Generalized Algorithm for Multi-Objective Reinforcement Learning and Policy Adaptation. In *Advances in Neural Information Processing Systems*; Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R., Eds.; Curran Associates, Inc.: 2019; pp. 14636–14647.