

BLACK BOX RECURSIVE TRANSLATIONS FOR MOLECULAR OPTIMIZATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Machine learning algorithms for generating molecular structures offer a promising new approach to drug discovery. We cast molecular optimization as a translation problem, where the goal is to map an input compound to a target compound with improved biochemical properties. Remarkably, we observe that when generated molecules are iteratively fed back into the translator, molecular compound attributes improve with each step. We show that this finding is invariant to the choice of translation model, making this a "black box" algorithm. We call this method Black Box Recursive Translation (BBRT), a new inference method for molecular property optimization. This simple, powerful technique operates strictly on the inputs and outputs of any translation model. We obtain new state-of-the-art results for molecular property optimization tasks using our simple drop-in replacement with well-known sequence and graph-based models. Our method provides a significant boost in performance relative to its non-recursive peers with just a simple "for" loop. Further, BBRT is highly interpretable, allowing users to map the evolution of newly discovered compounds from known starting points.

1 INTRODUCTION

Automated molecular design using generative models offers the promise of rapidly discovering new compounds with desirable properties. Chemical space is large, discrete, and unstructured, which together, present important challenges to the success of any molecular optimization campaign. Approximately 10^8 compounds have been synthesized (Kim et al., 2015) while the range of potential druggable candidates is estimated to be between 10^{23} and 10^{80} (Polishchuk et al., 2013). Consequently, new methods for intelligent search are paramount.

A recently introduced paradigm for compound generation treats molecular optimization as a translation task where the goal is to map an input compound to a target compound with favorable properties (Jin et al., 2019). This framework has presented impressive results for constrained molecular property optimization where generated compounds are restricted to be structurally similar to the source molecule.

We extend this framework to unconstrained molecular optimization by treating inference as a first-class citizen. We observe the remarkable fact that generated molecules can be repeatedly fed back into the model to generate even better compounds. This finding is invariant to the choice of translation model, making this a "black box" algorithm. This invariance is particularly attractive considering the recent emphasis on new molecular representations (Gómez-Bombarelli et al., 2018; Jin et al., 2018; Dai et al., 2018; Li et al., 2018; Kusner et al., 2017; Krenn et al., 2019). Using our simple drop-in replacement, our method can leverage these recently introduced molecular representations in a translation setting for better optimization.

We introduce Black Box Recursive Translation (BBRT), a new inference method for molecular property optimization. Surprisingly, by applying BBRT to well-known sequence and graph based models in the literature, we can produce new state-of-the-art results on property optimization benchmark tasks. Through an exhaustive exploration of various decoding strategies, we demonstrate the empirical benefits of using BBRT. We introduce simple ranking methods to decide which outputs are fed back into the model and find ranking to be an appealing approach to multi-property optimization. Finally, we demonstrate how BBRT is an extensible tool for interpretable and user-centric molecular design applications.

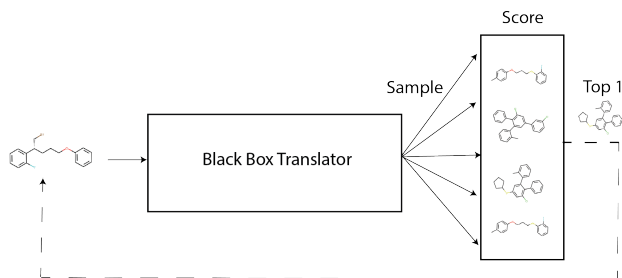


Figure 1: Black Box Recursive Translation (BBRT).

2 RELATED WORK

De novo molecular design. Recent work has focused on learning new molecular representations including graphs (You et al., 2018b; Li et al., 2018), grammars (Kusner et al., 2017; Dai et al., 2018), trees (Jin et al., 2018), and sequences (Gómez-Bombarelli et al., 2018; Krenn et al., 2019). Provided with a molecular representation, latent variable models (Kusner et al., 2017; Dai et al., 2018; Jin et al., 2018), markov chains (Seff et al., 2019), or auto-regressive models (You et al., 2018b) have been developed to learn distributions over molecular data. Molecular optimization has been approached with reinforcement learning (Popova et al., 2018; Zhou et al., 2019) and optimization in continuous, learned latent spaces (Gómez-Bombarelli et al., 2018).

We build on recent work introducing a third paradigm for design, focusing on molecular optimization as a translation problem (Jin et al., 2019). Jin and colleagues develop a graph based representation for molecular translation. Their method directly translates between a source graph and a target graph *via* an intermediate hidden representation. With many ways to modify a molecule to improve its biochemical properties, Jin et al. (2019) use stochastic hidden states coupled with a left-to-right greedy decoder to generate multi-modal outputs. While Jin et al. (2019) focus their efforts on constrained optimization, we focus on extending this framework to a more general unconstrained setting. We show that while their translation model is not fundamentally limited to constrained optimization, their inference method restricts the framework’s application to more general problems.

Matched molecular pair (MMP) analysis. MMP analysis is a popular cheminformatics framework for analyzing structure-property relationships (Turk et al., 2017). MMPs are extracted by mining databases of measured chemical properties and identifying pairs of chemical structures that share the same core and differ only by a small, well-defined structural difference, *e.g.*, where a methyl group is replaced by an isopropyl group (Tyrchan & Evertsson, 2017). The central hypothesis underlying MMPs is that the chemical properties of a series of closely related structures can be described by piecewise-independent contributions that various structural adducts make to the properties of the core.

MMP analysis has become a mainstream tool for interpreting the complex landscape of structure-activity relationships *via* simple, local perturbations that can be learnt and potentially transferred across drug design projects (Kubinyi, 1988). This framework serves as the chemistry analogue to a popular interpretability tool in machine learning called local interpretable model-agnostic explanations (LIME) (Ribeiro et al., 2016). MMP and LIME alike learn local surrogate models to explain individual predictions.

We view molecular translation as a learned MMP analysis. While Jin et al. (2019) use neural networks to learn a single high-dimensional MMP step, we extend this framework to infer a sequence of MMPs, extending the reach of this framework to problems beyond constrained optimization.

Translation models. Machine translation models composed of end-to-end neural networks (Sutskever et al., 2014) have enjoyed significant success as a general-purpose modeling tool for many applications including dialogue generation (Vinyals & Le, 2015) and image captioning (Vinyals et al., 2015). We focus on a recently introduced application of translation modeling, one of molecular optimization (Jin et al., 2019).

The standard approach to inference – approximately decoding the most likely sequence under the model – involves a left-to-right greedy search, which is known to be highly suboptimal, producing generic sequences exhibiting low diversity (Holtzman et al., 2019). Recent work propose diverse beam search (Li & Jurafsky, 2016; Vijayakumar et al., 2018; Kulikov et al., 2018), sampling meth-

ods geared towards open-ended tasks (Fan et al., 2018; Holtzman et al., 2019), and reinforcement learning for post-hoc auxiliary objective optimization (Wiseman et al., 2018; Shen et al., 2016; Bahdanau et al., 2017). Motivated by molecular optimization as a translation task, we develop Black Box Recursive Translation (BBRT). We show BBRT generates samples with better molecular properties than its non-recursive peers for both deterministic and stochastic decoding strategies.

3 MOLECULAR OPTIMIZATION AS A TRANSLATION PROBLEM.

For illustrative purposes, we describe the translation framework and the corresponding inference method for sequences. We emphasize that our inference method is a black box, which means it is invariant to specific architecture and representational choices.

Background. Our goal is to optimize the chemical properties of molecules using a sequence-based molecular representation. We are given $(x, y) \in (\mathcal{X}, \mathcal{Y})$ as a sequence pair, where $x = (x_1, x_2, \dots, x_m)$ is the source sequence with m tokens and $y = (y_1, y_2, \dots, y_n)$ is the target sequence with n tokens, and \mathcal{X} and \mathcal{Y} are the source and target domains respectively. We focus on problems where the source and target domains are identical. By construction, our training pairs (x, y) have high chemical similarity, which helps the model learn local edits to x . In addition, y is a paraphrase of x with improved property values. These properties are specified beforehand when constructing training pairs. A single task will typically optimize a single property such as potency, toxicity, or lipophilic efficiency.

A sequence to sequence (Seq2Seq) model learns parameters θ that estimate a conditional probability model $P(y|x; \theta)$, where θ is estimated by maximizing the log likelihood:

$$L(\theta) = \sum_{(x,y) \in (\mathcal{X}, \mathcal{Y})} \log P(y|x, \theta) \quad (1)$$

The conditional probability is typically factorized according to the chain rule: $P(y|x; \theta) = \prod_{t=1}^n P(y_t|y_{<t}, x, \theta)$. These models use an encoder-decoder architecture where the input and output are both parameterized by recurrent neural networks (RNNs). The encoder reads the source sequence x and generates a sequence of hidden representations. The decoder estimates the conditional probability of each target token given the source representations and its preceding tokens. The attention mechanism (Bahdanau et al., 2014) helps with token generation by focusing on token-specific source representations.

4 BLACK BOX RECURSIVE TRANSLATION

4.1 CURRENT INFERENCE METHODS

For translation models, the inference task is to compute $y^* = \arg \max p(y|x, \theta)$. Because the search space of potential outputs is incredibly large, in practice, we can only explore a limited number of sequences. Given a fixed computational budget, likely sequences are typically generated with heuristics. Decoding methods can be classified as deterministic or stochastic. We now describe both classes of decoding strategies in detail. For this section, we follow the notation described in Welleck et al. (2019).

Deterministic Decoding. Two popular deterministic methods include greedy search and beam search. Greedy search performs a single left-to-right pass, selecting the most likely token at each time step: $y_t = \arg \max p(y_t|y_{<t}, x, \theta)$. While this method is efficient, it leads to suboptimal generation as it does not take into account the future when selecting tokens.

Beam search is a generalization of greedy search and maintains a set of k hypotheses at each time-step where each hypothesis is a partially decoded sequence. While in machine translation beam search variants are the preferred method, for more open-ended tasks, beam search fails to generate diverse candidates. Recent work has explored diverse beam search (Li & Jurafsky, 2016; Vijayakumar et al., 2018; Kulikov et al., 2018). These methods address the reduction of number of duplicate sequences to varying extents, thereby increasing the entropy of generated sequences.

Stochastic Decoding. A separate class of decoding methods sample from the model at generation time, $y_t \sim q(y_t|y_{<t}, x, p_\theta)$. This method has shown to be effective at generating diverse samples and can better explore target design spaces. We consider a top- k sampler (Fan et al., 2018),

which restricts sampling to the k most-probable tokens at time-step t . This corresponds to restricting sampling to a subset of the vocabulary $U \subset V$. U is the subset of V that maximizes $\sum_{y \in U} p_\theta(y_t | y_{y < t}, x)$:

$$q(y_t | y_{<t}, x, p_\theta) = \begin{cases} \frac{p_\theta(y_t | y_{<t}, x)}{Z} & y_t \in U \\ 0 & \text{otherwise} \end{cases}$$

4.2 RECURSIVE INFERENCE

We are given $(x, y) \in (X, Y)$ as a sequence pair where by construction (x, y) have high chemical similarity and y is a paraphrase of x with an improved value for a prespecified property. At test time, we are interested in recursively inferring new sequences. Let y_i denote a random sequence for recursive iteration i . Let $\{y_i^{(k)}\}_{k=1}^K$ be a set of K outputs generated from $p_\theta(y_i | x)$ when $i = 0$. We use a scoring function S to compute the best of K outputs denoted as \hat{y}_i . For $i > 0$, we infer K outputs from $p_\theta(y_i | \hat{y}_{i-1})$. This process is repeated for n iterations.

Scoring functions. In principle, all K outputs at iteration i can be independently conditioned on to generate new candidates for iteration $i + 1$. This procedure scales exponentially with respect to space and time as a function of n iterations. Therefore, we introduce a suite of simple ranking strategies to score K output sequences to decide which output becomes the next iteration’s source sequence. We consider a likelihood based ranking as well as several chemistry-specific metrics further described in the experiments. Designing well-informed scoring functions can help calibrate the distributional statistics of generated sequences, aid in multi-property optimization, and provide interpretable sequences of iteratively optimal translations.

Ensemble outputs. After n recursive iterations, we ensemble the generated outputs $\{y_0, y_1, \dots, y_n\}_{k=1}^K$ and score the sequences on a desired objective. For property optimization, we return the arg max. In principle, BBRT is not limited to ensembling recursive outputs from a *single* model. Different modeling choices and molecular representations have different inductive biases, which influence the diversity of generated compounds. BBRT can capitalize on these differences by providing a coherent method to aggregate results.

5 EXPERIMENTS

We apply BBRT to solve unconstrained and multi-property optimization problems. To highlight the generality of our method, we apply recursive translation to both sequence and graph-based translation models. **Surprisingly, we show that BBRT generates state-of-the-art results on molecular property optimization using already published modeling approaches.** Next we describe how recursive inference lends itself to interpretability through the generation of molecular traces, allowing practitioners to map the evolution of discovered compounds from known starting points through a sequence of local structural changes. At any point in a molecular trace, users may introduce a "break point" to consider alternative translations thereby personally evaluating the tradeoffs between conflicting design objectives. Finally, we apply BBRT to multi-property optimization.

Models. We apply our inference method to sequence and graph-based molecular representations. For sequences, we use the recently introduced SELFIES molecular representation (Krenn et al., 2019), a sequence based representation for semantically constrained graphs. Empirically, this method generates a high percentage of valid compounds (Krenn et al., 2019). Using SELFIES, we develop a sequence to sequence model with an encoder-decoder framework. The encoder and decoder are both parameterized by RNNs with LSTM cells. The encoder is a 2-layer bidirectional RNN and the decoder is a 1-layer unidirectional forward RNN. We also use attention (Bahdanau et al., 2014) for decoding. The hidden representations are non-probabilistic and are optimized to minimize a standard cross-entropy loss with teacher forcing. Decoding is performed using deterministic and stochastic decoding strategies described in Section 4.1. For graphs, we use a tree-based molecular representation (Jin et al., 2018) with the exact architecture described in (Jin et al., 2019). Decoding is performed using a probabilistic extension with latent variables described in (Jin et al., 2019)—we sample from the prior k times followed by left to right greedy decoding.

Method	Penalized logP			QED		
	1st	2nd	3rd	1st	2nd	3rd
ZINC	4.52	4.30	4.23	0.948	0.948	0.948
ORGAN	3.63	3.49	3.44	0.896	0.824	0.820
JT-VAE	5.30	4.93	4.49	0.925	0.911	0.910
GCPN	7.98	7.85	7.80	0.948	0.947	0.946
VJTNN (6600 samples)	5.97	4.96	4.71	0.948	0.948	0.948
Seq2Seq (6600 samples)	4.65	4.53	4.49	0.948	0.948	0.948
BBRT+VJTNN	10.13	10.10	9.91	0.948	0.948	0.948
BBRT+Seq2Seq	6.74	6.47	6.42	0.948	0.948	0.948

Table 1: Top 3 property scores on penalized logP and QED tasks.

Data. We construct training data by sampling molecular pairs (X, Y) with molecular similarity $sim(X, Y) > \tau$ and property improvement $\delta(Y) > \delta(X)$ for a given property δ . Constructing training pairs with a similarity constraint can help avoid degenerate mappings. In contrast to Jin et al. (2019), we only enforce the similarity constraint for the construction of training pairs. Similarity constraints are not enforced at test time. Molecular similarity is measured by computing Tanimoto similarity with Morgan fingerprints Rogers & Hahn (2010). All models were trained on the open-source ZINC dataset.

Properties. For all experiments, we focus on optimizing two well-known drug properties of molecules. First, we optimize the water-octanol partition coefficient (logP). Similar to (Jin et al. (2018); Kusner et al. (2017); Gómez-Bombarelli et al. (2018)), we consider a penalized logP score that incorporates ring size and synthetic accessibility. Following Jin et al. (2019) we extracted 99K translation pairs for training using a similarity constraint of 0.4. Second, we optimize quantitative estimate of drug likeness (QED) Bickerton et al. (2012). Following Jin et al. (2019), we construct training pairs where the source compound has a QED score within the range [.7 .8] and the target compound has a QED score within the range [.9 1.0]. While Jin et al. (2019) evaluates QED performance based on a closed set translation task, we evaluate this model for unconstrained optimization. We extracted a training set of 88K molecule pairs. Please see the supplement for details on how these properties were computed.

Scoring functions. Here we describe scoring functions that are used to rank K outputs for recursive iteration i . The highest scoring sequence is used as the next iteration’s source compound..

- Penalized LogP: Choose the compound with the max LogP value. This is useful when optimizing LogP as a primary or auxiliary property.
- QED: Choose the compound with the max QED value. This is useful when optimizing QED as a primary or auxiliary property.
- Max Delta Sim: Choose the compound with the highest chemical similarity to the previous iteration’s source compound. This is useful for interpretable, molecular traces by creating a sequence of translations with local edits.
- Max Init Sim: Choose the compound with the highest similarity to the initial seed compound. This is useful for input-constrained optimization.
- Min Mol Wt: Choose the compound with the minimum molecular weight. This is useful for rectifying a molecular generation artifact where models optimize logP by simply adding functional groups, thus increasing the molecular weight (see supplement fig.).

Diversity is computed as follows. Let Y be a set of translated compounds. Consistent with the literature we report diversity as

$$DIV(Y) = \frac{1}{|Y|(|Y| - 1)} \sum_{y \in Y} \sum_{y' \in Y, y'} 1 - \delta(y, y') \quad (2)$$

δ is the Tanimoto similarity computed on the Morgan fingerprints of y and y' .




Figure 2: Left and Center: Top 100 logP generated compounds under BBRT applied to Seq2Seq and Graph2Graph and their non-recursive counterparts. Right: Diversity of top 100 generated compounds under both BBRT models and the top 100 compounds from the training data.




Figure 3: Top scoring compounds for properties logP and QED under BBRT+Seq2Seq and BBRT+Graph2Graph.

5.1 MOLECULE GENERATION RESULTS

Setup. For property optimization, the goal is to generate compounds with the highest possible penalized logP and QED values. For notation, we denote BBRT applied to model X as R-X. We consider R-Graph2Graph and R-Seq2Seq under 3 decoding strategies and 5 scoring functions. For the logP task, we seed our translations with 900 maximally diverse compounds with an average pairwise diversity of 0.94 ± 0.04 relative to $.86 \pm .04$, which is the average pairwise diversity of the training data. Maximally diverse compounds were computed using the MaxMin algorithm (Ashton et al., 2002) on Morgan fingerprints. We found seeding our translations with diverse compounds helped BBRT generate higher scoring compounds. For both BBRT applications, we sample 100 times from a top-2 and from a top-5 sampler and aggregate these outputs with a beam search with 20 beams outputting 20 compounds.

Baselines. We compare our method with the following state-of-the-art baselines. Junction Tree VAE (Jin et al., 2018) combines a graph-based representation with latent variables for molecular graph generation. Molecular optimization is performed using Bayesian Optimization on the learned latent space to search for compounds with optimized property scores. JT-VAE has been shown to outperform other methods including Character-VAE (Gómez-Bombarelli et al., 2018), Grammar-VAE (Kusner et al., 2017), and Syntax-Directed-VAE (Dai et al., 2018). We also compare against two reinforcement learning molecular generation methods; ORGAN (Guimaraes et al., 2017) uses SMILES strings (Weininger, 1988), a text-based molecular representation, and GCPN (You et al., 2018a), which uses graphs.

We also compare against non-recursive variants of the Seq2Seq and Graph2Graph models considered here. For a fair comparison, we admit similar computational budgets to these baselines. For Seq2Seq we include a deterministic and stochastic decoding baseline. For the deterministic baseline, we use beam search with 20 beams and compute the 20 most probable sequences under the model. For the stochastic baseline, we sample 6600 times from a top-5 sampler. For details on the sampler, we refer readers to section 4.1. For VJTNN (Jin et al., 2019), the Graph2Graph approach, we use their reported Github implementation. There is not an obvious stochastic and deterministic parallel

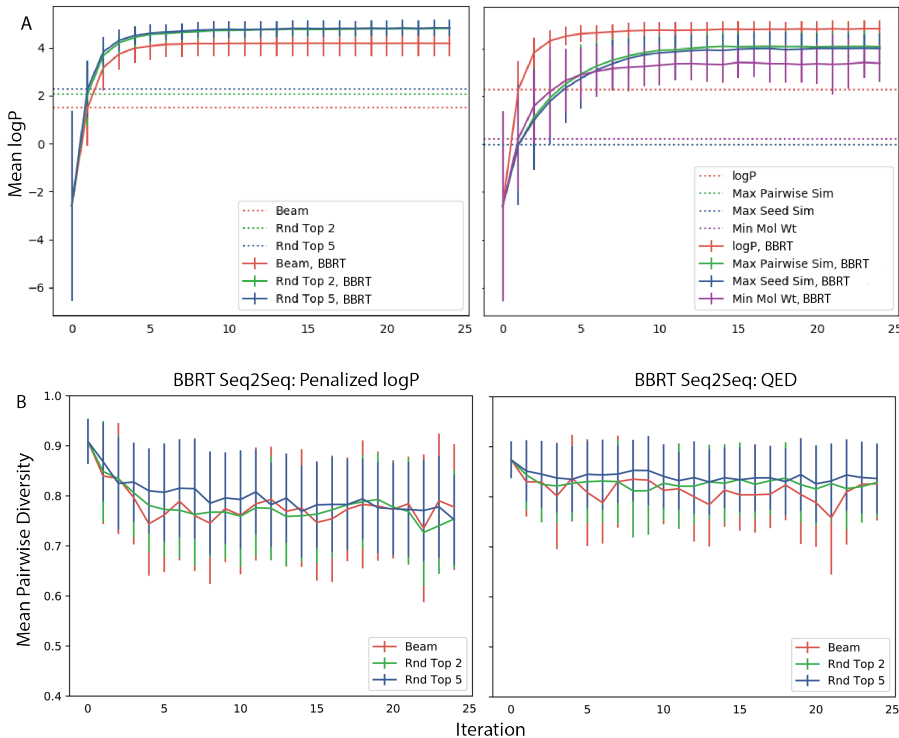


Figure 4: Ablation experiments using BBRT applied to Seq2Seq. A. Left: Mean logP from 900 translations as a function of recursive iteration for three decoding strategies. Dotted lines denote non-recursive counterparts. Right: Mean logP as a function of recursive iteration for 4 scoring functions. B. Left: Diversity of generated outputs across recursive iterations for logP translation. Right: Diversity of generated outputs across recursive iterations for QED translation.

considering their method is probabilistic. Therefore, we focus on comparing to a fair computational budget by sampling 6600 times from the prior distribution followed by greedy decoding. For fair comparisons to the recursive application, the same corresponding sampling strategies are used, with 100 samples per iteration.

Results. Following reporting practices in the literature, we report the best 3 property scores found by each model and the fraction of generated molecules that satisfy chemical validity. Table 1 summarizes these results. The best 3 property scores found in ZINC250k are included for comparison. For logP optimization, BBRT+VJTNN (R-Graph2Graph) significantly outperforms all baseline models including VJTNN (Graph2Graph), Seq2Seq, and BBRT+Seq2Seq (R-Seq2Seq). R-Seq2Seq outperforms Seq2Seq, highlighting the benefits of recursive inference for both molecular representations. For QED property optimization, the two translation models and the BBRT variants all find the same top 3 property scores, which is a new state-of-the-art result for QED optimization.

In Fig. 2, we report the top 100 logP compounds generated by both BBRT applications relative to its non-recursive counterparts and observe significant improvements in logP from using BBRT. We also report a diversity measure of the generated candidates for both BBRT models and the top 100 logP compounds in the training data. The Graph2Graph variant produces logP compounds that are more diverse than the compounds in the training data, while the compounds generated by Seq2Seq are less diverse.

Fig. 3 visualizes the top 2 discovered compounds by R-Graph2Graph and R-Seq2Seq under both properties. For logP, while R-Graph2Graph produces compounds with higher property values, R-Seq2Seq’s top 2 generated compounds have a richer molecular vocabulary. R-Seq2Seq generates compounds with heterocycles and linkages while R-Graph2Graph generates a chain of linear hydrocarbons, which looks quite similar to the top reported compounds in GCPN You et al. (2018b), an alternative graph-based representation. This distinction in the vocabulary highlights the usefulness of a sequence-based representation and a flexible framework to aggregate results across molecular representations.

Differences between logP and QED. For logP, BBRT provides a 27% improvement over state-of-the-art for property optimization, while for QED, despite recursive methods outperforming the baselines, our non-recursive Seq2Seq and Graph2Graph baselines perform just as well. We note this is not already reported in the literature. This highlights an important difference between explorative and interpolative design. LogP optimization is an explorative task—the property is unbounded and generative models have produced better values than what is in the dataset. While QED is an interpolative task—the property is bounded and the best reported compound values exist in the dataset.

5.2 EMPIRICAL PROPERTIES OF RECURSIVE TRANSLATION.

We perform a sequence of ablation experiments to better understand the effect of various BBRT design choices on performance. We highlight the variability in average logP from translated outputs at each iteration with different decoding strategies (Fig. 4A left) and scoring functions (Fig. 4A right).

On the importance of stochastic decoding. For non-recursive and recursive translation models, **stochastic decoding methods outperformed deterministic methods** on average logP scores (Fig. 4A left) and average pairwise diversity (Fig. 4B) for generated compounds as a function of recursive iteration. Non-greedy-based search strategies are not common practice in *de novo* molecular design (Gómez-Bombarelli et al., 2018; Kusner et al., 2017; Jin et al., 2019). While recent work emphasizes novel network architectures and generating diverse compounds using latent variables (Gómez-Bombarelli et al., 2018; Kusner et al., 2017; Jin et al., 2018), we identify an important design choice that typically has been under-emphasized. This trend has also been observed in the NLP community where only recently have researchers highlighted the importance of well-informed search techniques (Kulikov et al., 2018).

Regardless of the decoding strategy, we observed improvements in mean logP with iterations (Fig. 2A) when using BBRT. When optimizing for logP, a logP scoring function quickly discovers the best scoring compounds while auxiliary scoring functions improve logP at a slower rate and do not converge to the same scores (2A right). This tradeoff highlights the role of conflicting molecular design objectives.

For Fig. 4A, the standard deviation typically decreased with iteration number. Property values concentrate to a certain range. With property values concentrating, it is reasonable to question whether BBRT produces compounds with less diversity. In Fig. 4B we show average pairwise diversity of translated outputs per recursive iteration across 3 decoding strategies and observe decay in diversity for logP. For the best performing decoding strategy, the top-5 sampler, diversity decays from about 0.86 after a single translation to about 0.78 after 25 recursive translations. This decay is somewhat a product of the data—higher logP values tend to be less diverse than a random set of compounds. For QED (Fig. 4B right), we observe limited decay. Differences in decay rate might be attributed to task variability, one being explorative and the other interpolative.

5.3 INTERPRETABLE, USER CENTRIC OPTIMIZATION

Generative models can produce a large number of novel, promising compounds. For active drug-design projects, the sheer number of suggestions produces a practical challenge. Project teams operate within limited synthesis budgets and must prioritize a few suggestions to synthesize and experimentally validate. From observing medicinal chemists work with molecular generative models for active projects, we found that chemists rank generated compounds based on intuiting how a model might have arrived at that suggestion. Therefore, providing interpretable paths of optimization can help drive user adoption.

Molecular traces. BBRT generates a sequence of iteratively optimal local edits from an initial compound to an optimized one. We call optimization paths based on local structural changes a *molecular trace* (Fig. 5A). In Gómez-Bombarelli et al. (2018), latent space interpolation methods are used to show the intermediate sequence of compounds between an initial starting point and an optimized compound. Their method, however, had a low validity rate for new compound generation and did not present a method to allow users to reconsider steps in the optimization path based on their input. RL based molecular generative techniques have not emphasized interpretable molecular modeling. Recent RL methods focus on molecular construction and therefore are not well-suited for the generation of molecular traces, which requires a model to move *between* compounds. Here we




Figure 5: A. Generated molecular trace by ranking intermediate outputs by the maximum pairwise Tanimoto (TM) similarity. B. An example molecular breakpoint. Alternative translations are considered from compound (2) each with its own design tradeoffs.

develop molecular traces for translation-based design. We use the Min Delta Sim scoring function to generate traces that have the minimum possible structural changes while still improving logP.

Two features of molecular traces make them particularly well-suited for pharmaceutical research:

- This idea of iteratively optimal *and interpretable* local edits has a long history in pharmaceutical research. The Free-Wilson approach and its variants build local models for structure-activity relationships using molecular subgraphs themselves as model descriptors, with the biological activity of molecules being described by linear summations of activity contributions of specific subgraphs (Eriksson et al., 2014). These approaches provide simple and straightforward interpretability within the chemical space spanned by the subgraphs, and molecular traces are equivalent to a sequence of Free-Wilson steps towards optimal molecules.
- Drug design projects proceed *via* synthesis-test-analysis/ideation cycles, and since projects stay away from well-explored chemical spaces for intellectual property reasons, the synthesis of compounds proposed during each cycle almost always requires chemists to come up with new synthetic routes. Synthetic chemists formulate synthesis routes as multi-step combinations of a small number of well-known chemical reactions that carry out local transformations to molecules until the desired structure is obtained. Since molecular traces are themselves local transformations of structures from a seed compound to an optimal compound, synthetic chemists can carry out the individual steps of a molecular trace to synthesize the optimal compound generated by BBRT.

Molecular breakpoints. Molecular design requires choosing between conflicting objectives. Increased logP is correlated with poor oral drug-like properties and rapid clearance from the body Ryckmans et al. (2009). Increasing the QED score might translate to a compound that is structurally dissimilar from the seed compound, which could result in an inactive compound against a protein of interest. Building on molecular traces, our method allows users to “debug” any step of the translation process and consider alternative steps, just like breakpoints in computer programs. Given the difficulty of specifying a single global objective function for molecular design, methods that can allow for chemists to be a part of the design loop are critical for user adoption and sensible design. In Fig. 5B, we show an example from an R-Seq2Seq model trained to optimize logP. Here we revisit the translation from step (2) to step (3) in Fig. 5A by considering four alternatives picked from 100 stochastically decoded compounds. These alternatives require evaluating the tradeoffs between logp, QED, molecular weight, the number of rotational bonds, and chemical similarity with compound (2).

5.4 MULTI-PROPERTY OPTIMIZATION

We consider auxiliary property optimization by ranking recursive outputs using the scoring function. This function decides what compound is propagated to the next recursive step. We apply BBRT to Seq2Seq modeling (R-Seq2Seq) and use the trained QED translator described in Section 5.1. The inference task is to optimize QED and logP as the primary and secondary properties respectively. We compare scoring outputs based on QED and logP. In Fig. 6A, we compute the average logP

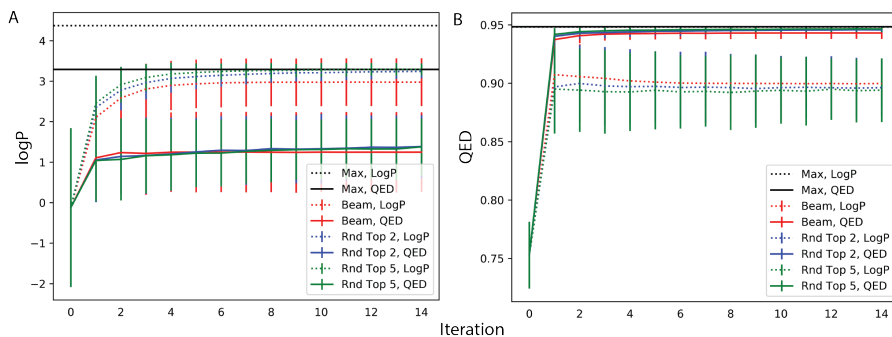


Figure 6: Applying BBRT to multi-property optimization. QED is the primary target and logP is the auxiliary property. A: Average logP as a function of recursive iteration for three decoding strategies with primary and auxiliary property scoring functions. B: Average QED as a function of recursive iteration for three decoding strategies with primary and auxiliary property scoring functions.

per recursive iteration for a set of translated compounds across three decoding strategies. Dotted lines optimize logP as the scoring function while the solid lines optimize QED. For both scoring functions, we report the max logP value generated. For all decoding strategies, average logP reaches higher values under scoring by logP relative to scoring by QED. In Fig. 6B, we plot average QED values using the same setup and observe that optimizing logP still significantly improves the QED values of generated compounds. This method also discovers the same max QED value as scoring by QED. This improvement, however, has tradeoffs in the limit for average QED values generated. After 15 recursive iterations the *average* QED values of the generated compounds under a logP scoring function converge to lower values relative to QED values for compounds scored by QED for all three decoding strategies. We repeat this experiment with Graph2Graph and show similar effects in the supplement (Fig. 8).

Multi-property optimization by ranking extends to variables that are at minimum loosely positively correlated. For QED optimization, the average logP value for unoptimized QED compounds is 0.30 ± 1.96 while for optimized QED compounds the average logP value is 0.79 ± 1.45 . Additionally, QED compounds in the target range [9 1.0] in the training data had a positive correlation between its logP and QED values (Spearman rank correlation; $\rho = 0.07$ $P < 0.026$). This correlation did not hold for QED compounds in the range [7 .8] unoptimized QED compounds ($\rho = .007$, $P < .8$).

6 FUTURE WORK

We develop BBRT for molecular optimization. BBRT is a simple algorithm that feeds the output of translation models back into the same model for additional optimization. We apply BBRT to well-known models in the literature and produce new state-of-the-art results for property optimization tasks. We describe molecular traces and user centric optimization with molecular breakpoints. Finally, we show how BBRT can be used for multi-property optimization. For future work, we will extend BBRT to consider multiple translation paths simultaneously. In addition we will extend translation models to low-resource settings, where property annotations are expensive to collect.

ACKNOWLEDGMENTS

REFERENCES

- Mark Ashton, John Barnard, Florence Casset, Michael Charlton, Geoffrey Downs, Dominique Gorse, John Holliday, Roger Lahana, and Peter Willett. Identification of diverse database subsets using Property-Based and Fragment-Based molecular descriptions. *Quant. Struct.-Act.Relat.*, 21 (6):598–604, December 2002.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural Machine Translation by Jointly Learning to Align and Translate. *arXiv.org*, September 2014.
- Dzmitry Bahdanau, Philemon Brakel, Kelvin Xu, Anirudh Goyal, Ryan Lowe, Joelle Pineau, Aaron Courville, and Yoshua Bengio. An Actor-Critic Algorithm for Sequence Prediction. In *International Conference on Learning Representations*, 2017.

- G Richard Bickerton, Gaia V Paolini, Jérémy Besnard, Sorel Muresan, and Andrew L Hopkins. Quantifying the chemical beauty of drugs. *Nature chemistry*, 4(2):90, 2012.
- Hanjun Dai, Yingtao Tian, Bo Dai, Steven Skiena, and Le Song. Syntax-directed variational autoencoder for structured data. In *International Conference on Learning Representations*, 2018.
- Mats Eriksson, Hongming Chen, Lars Carlsson, J. Willem M. Nissink, John G. Cumming, and Ingemar Nilsson. Beyond the Scope of Free-Wilson Analysis. 2: Can Distance Encoded R-Group Fingerprints Provide Interpretable Nonlinear Models? *Journal of Chemical Information and Modeling*, 54(4):1117–1128, 2014. ISSN 1549-9596. doi: 10.1021/ci500075q.
- Angela Fan, Mike Lewis, and Yann N Dauphin. Hierarchical Neural Story Generation. *ACL*, pp. 889–898, 2018.
- Rafael Gómez-Bombarelli, Jennifer N Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D Hirzel, Ryan P Adams, and Alán Aspuru-Guzik. Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules. *ACS Central Science*, 4(2):268–276, January 2018.
- Gabriel Lima Guimaraes, Benjamín Sánchez-Lengeling, Carlos Outeiral, Pedro Luis Cunha Farias, and Alán Aspuru-Guzik. Objective-Reinforced Generative Adversarial Networks (ORGAN) for Sequence Generation Models. *arXiv.org*, May 2017.
- Ari Holtzman, Jan Buys, Maxwell Forbes, and Yejin Choi. The Curious Case of Neural Text Degeneration. *arXiv.org*, April 2019.
- Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Junction tree variational autoencoder for molecular graph generation. In *International Conference on Machine Learning*, pp. 2328–2337, 2018.
- Wengong Jin, Kevin Yang, Regina Barzilay, and Tommi Jaakkola. Learning multimodal graph-to-graph translation for molecule optimization. In *International Conference on Learning Representations*, 2019.
- Sunghwan Kim, Paul A Thiessen, Evan E Bolton, Jie Chen, Gang Fu, Asta Gindulyte, Liany Han, Jane He, Siqian He, Benjamin A Shoemaker, et al. Pubchem substance and compound databases. *Nucleic acids research*, 44(D1):D1202–D1213, 2015.
- Mario Krenn, Florian Häse, AkshatKumar Nigam, Pascal Friederich, and Alán Aspuru-Guzik. SELFIES: a robust representation of semantically constrained graphs with an example application in chemistry. *arXiv.org*, May 2019.
- Hugo Kubinyi. Free wilson analysis. theory, applications and its relationship to hansch analysis. *Quantitative Structure-Activity Relationships*, 7(3):121–133, 1988.
- Ilya Kulikov, Alexander H Miller, Kyunghyun Cho, and Jason Weston. Importance of a Search Strategy in Neural Dialogue Modelling. *arXiv.org*, November 2018.
- Matt J Kusner, Brooks Paige, and José Miguel Hernández-Lobato. Grammar variational autoencoder. In *International Conference on Machine Learning*, pp. 1945–1954, 2017.
- Jiwei Li and Dan Jurafsky. Mutual Information and Diverse Decoding Improve Neural Machine Translation. *arXiv.org*, January 2016.
- Yujia Li, Oriol Vinyals, Chris Dyer, Razvan Pascanu, and Peter Battaglia. Learning Deep Generative Models of Graphs. *arXiv.org*, March 2018.
- Pavel G Polishchuk, Timur I Madzhidov, and Alexandre Varnek. Estimation of the size of drug-like chemical space based on gdb-17 data. *Journal of computer-aided molecular design*, 27(8): 675–679, 2013.
- Mariya Popova, Olexandr Isayev, and Alexander Tropsha. Deep reinforcement learning for de novo drug design. *Science Advances*, 4(7):eaap7885, July 2018.

- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. Why should i trust you?: Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD 16*, 2016.
- David Rogers and Mathew Hahn. Extended-connectivity fingerprints. *Journal of chemical information and modeling*, 50(5):742–754, 2010.
- Thomas Ryckmans, Martin P. Edwards, Val A. Horne, Ana Monica Correia, Dafydd R. Owen, Lisa R. Thompson, Isabelle Tran, Michelle F. Tutt, and Tim Young. Rapid assessment of a novel series of selective CB2 agonists using parallel synthesis protocols: A Lipophilic Efficiency (LipE) analysis. *Bioorganic and Medicinal Chemistry Letters*, 2009. ISSN 0960894X. doi: 10.1016/j.bmcl.2009.05.062.
- Ari Seff, Wenda Zhou, Farhan Damani, Abigail Doyle, and Ryan P Adams. Discrete Object Generation with Reversible Inductive Construction. *arXiv.org*, July 2019.
- Shiqi Shen, Yong Cheng, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. Minimum risk training for neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1683–1692, 2016.
- Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. Sequence to sequence learning with neural networks. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, NIPS’14, pp. 3104–3112, Cambridge, MA, USA, 2014. MIT Press.
- Samo Turk, Benjamin Merget, Friedrich Rippmann, and Simone Fulle. Coupling Matched Molecular Pairs with Machine Learning for Virtual Compound Optimization. *Journal of Chemical Information and Modeling*, 57(12):3079–3085, dec 2017. ISSN 1549-9596. doi: 10.1021/acs.jcim.7b00298.
- Christian Tyrchan and Emma Evertsson. Matched Molecular Pair Analysis in Short: Algorithms, Applications and Limitations. *Computational and Structural Biotechnology Journal*, 15:86–90, jan 2017. ISSN 20010370. doi: 10.1016/j.csbj.2016.12.003.
- Ashwin K Vijayakumar, Michael Cogswell, Ramprasaath R Selvaraju, Qing Sun, Stefan Lee, David Crandall, and Dhruv Batra. Diverse Beam Search for Improved Description of Complex Scenes. *Thirty-Second AAAI Conference on Artificial Intelligence*, April 2018.
- Oriol Vinyals and Quoc Le. A Neural Conversational Model. *arXiv.org*, June 2015.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3156–3164, 2015.
- David Weininger. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of chemical information and computer sciences*, 28(1):31–36, 1988.
- Sean Welleck, Kianté Brantley, Hal Daumé Iii, and Kyunghyun Cho. Non-monotonic sequential text generation. In *International Conference on Machine Learning*, pp. 6716–6726, 2019.
- Sam Wiseman, Stuart Shieber, and Alexander Rush. Learning neural templates for text generation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 3174–3187, 2018.
- Jiaxuan You, Bowen Liu, Zhitao Ying, Vijay Pande, and Jure Leskovec. Graph convolutional policy network for goal-directed molecular graph generation. In *Advances in Neural Information Processing Systems*, pp. 6410–6421, 2018a.
- Jiaxuan You, Rex Ying, Xiang Ren, William Hamilton, and Jure Leskovec. Graphrnn: Generating realistic graphs with deep auto-regressive models. In *International Conference on Machine Learning*, pp. 5694–5703, 2018b.
- Zhenpeng Zhou, Steven Kearnes, Li Li, Richard N Zare, and Patrick Riley. Optimization of Molecules via Deep Reinforcement Learning. *Scientific reports*, 9(1):10752, July 2019.

A RECURSIVE PENALIZED LOGP EXPERIMENTS

Training details. For the Seq2Seq model, the hidden state dimension is 500. We use a 2 layer bidirectional RNN encoder and 1 layer unidirectional decoder with attention Bahdanau et al. (2017). The model was trained using an Adam optimizer for 20 epochs with learning rate 0.001. For the graph-based model, we used Jin et al. (2019)’s implementation, which can be downloaded from <https://github.com/wengong-jin/iclr19-graph2graph>.

Property calculation. Penalized logP is calculated using You et al. (2018a)’s implementation. Their implementation utilizes RDKit to compute clogP and synthetic accessibility scores. QED scores are also computed using RDKit.




Figure 7: Comparison of scoring functions for BBRT applied to Graph2Graph (R-Graph2Graph). Y-axis is mean logP from 900 translations as a function of recursive iteration. Dotted lines denote non-recursive counterparts. Rec. Inf. is synonymous with BBRT.




Figure 8: Applying BBRT to multi-property optimization with Graph2Graph model. QED is the primary target and logP is the auxiliary property. Left: Average logP as a function of recursive iteration under two scoring functions—QED and logP. Right: Average QED as a function of recursive iteration for same two scoring functions.