Black Box Recursive Translations for Molecular Optimization

Farhan Damani
Dept. of Computer Science
Princeton University
fdamani@princeton.edu

Vishnu Sresht
Simulation & Modeling Sciences
Pfizer Research & Development
vishnu.sresht@pfizer.com

Stephen Ra
Simulation & Modeling Sciences
Pfizer Research & Development
stephen.ra@pfizer.com

Abstract

Machine learning algorithms for generating molecular structures offer a promising 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. We observe that when generated molecules are reiteratively fed back into the translator, compound attributes improve with each step. We demonstrate 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 results competitive with state-of-the-art performance on 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. Lastly, BBRT is 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. A recent paradigm 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., 2019b]. We extend this framework to unconstrained molecular optimization by treating inference as a first-class citizen. We find that generated molecules can be recursively fed back into the model to generate compounds with improved property values. Moreover, we note an apparent invariance to the choice of translation model and argue this finding is relevant considering the 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 method, Black Box Recursive Translation (BBRT) (Figure 1), we leverage different sequence- and graph-based models from the literature for 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 secondary property optimization. Finally, we demonstrate how BBRT is an extensible tool for interpretable and user-centric molecular design applications.

Machine Learning for Molecules Workshop at NeurIPS 2020. https://ml4molecules.github.io

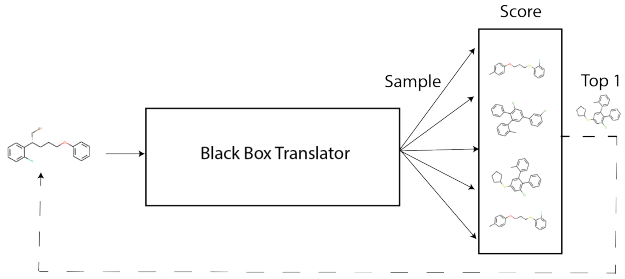

Figure 1: Black Box Recursive Translation (BBRT).

2 Framework

2.1 Model

As a baseline, we use a sequence-to-sequence (Seq2Seq) model [Sutskever et al., 2014] which learns parameters θ that estimate a conditional probability model $P(y|x;\theta)$, where θ is estimated by maximizing the log-likelihood:

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

The conditional probability is typically factorized according to the chain rule: $P(y|x;\theta) = \prod_{t=1}^{n} P(y_t|y_{<t},x,\theta)$. Our Seq2Seq model uses an encoder-decoder architecture, where the encoder and decoder are both parameterized by recurrent neural networks (RNNs) with Long Short-Term Memory (LSTM) cells [Hochreiter and Schmidhuber, 1997], and 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 both a deterministic strategy (beam search [Graves, 2012, Boulanger-Lewandowski et al., 2013]) and a stochastic decoding strategy using 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 $\sum_{y\in U} p_\theta(y_t|y_{y<t},x)$:

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

2.2 Recursive translation

For translation models, the inference task is to compute $y^* = \arg\max p(y|x,\theta)$. Given $(x,y) \in (X,Y)$ as a sequence pair where by construction (x,y) has high chemical similarity and y scores higher on a prespecified property compared to x, we construct training data using the ZINC dataset [Irwin et al., 2012] by sampling molecular pairs (X,Y) with Tanimoto similarity $sim(X,Y) > \tau$ and property improvement $\delta(Y) > \delta(X)$ for a given property δ. In contrast to Jin et al. [2019b], we only enforce the similarity constraint for the construction of training pairs.

At test time, we are interested in recursively inferring new sequences. Let y_i denote a random sequence for recursive iteration i and let $\{y_i^{(k)}\}_{k=1}^K$ be a set of K outputs generated from $p_\theta(y_i|x)$ when $i = 0$. To prioritize sequences for the next iteration, 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 $p_\theta(y_i|\hat{y}_{i-1})$. After n recursive iterations, we ensemble the generated outputs $\{y_0, y_1, \ldots, y_n\}_{k=1}^K$ and score the sequences on a desired objective. For property optimization, we return the $\arg\max$.

3 Results

To highlight the generality of our method, we apply recursive translation to established sequence and graph-based translation models [1] and compare against a number of established models in the literature including the Objective-Reinforced Generative Adversarial Network (ORGAN; Guimaraes et al. 2017), Junction Tree Variational Autoencoder (JT-VAE; Jin et al. 2018), Graph Convolutional

[1] We denote BBRT applied to model X as *'BBRT-X'*

	Penalized logP			QED		
Method	1st	2nd	3rd	1st	2nd	3rd
ZINC-250K	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
JTNN	5.97	4.96	4.71	0.948	0.948	0.948
Seq2Seq	4.65	4.53	4.49	0.948	0.948	0.948
BBRT-JTNN	**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.

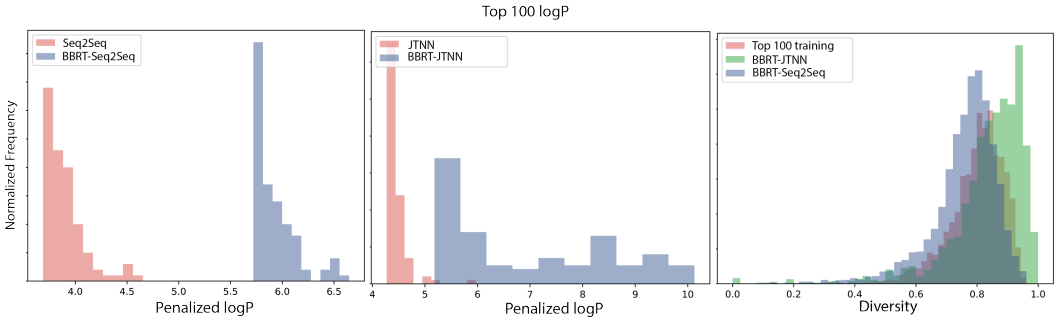

Figure 2: (Left and Center): Top 100 logP generated compounds under BBRT-Seq2Seq, BBRT-JTNN, and their non-recursive counterparts. (Right): Diversity of top 100 generated compounds under both BBRT models and the top 100 compounds from the training data.

Policy Network (GCPN; You et al. 2018), Variational Junction-Tree Encoder-Decoder (JTNN; Jin et al. 2019b), and Seq2Seq. We report the top three scores (penalized logP and QED) for each model (Table 1). We find that for logP optimization, BBRT-JTNN significantly outperforms all baseline models including the JTNN, Seq2Seq, and BBRT-Seq2Seq. BBRT-Seq2Seq outperforms Seq2Seq, highlighting the benefits of recursive inference for both molecular representations. For QED optimization, the two translation models and BBRT variants all find the same top three property scores.

In Figure 2, we report the top 100 logP compounds generated by both BBRT applications relative to their non-recursive counterparts and observe significant improvements in logP from using BBRT. Consistent with [Jin et al., 2019a], we also report diversity as $\text{DIV}(Y) = \frac{1}{|Y|(|Y|-1)} \sum_{y \in Y} \sum_{y' \in Y, y'} 1 - \delta(y, y')$ of the generated candidates for both BBRT models and the top 100 logP compounds in the training data. We find BBRT-JTNN produces logP compounds that are more diverse than the compounds in the training data, while the compounds generated by BBRT-Seq2Seq are less diverse.

3.1 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 (Figure 3A left) and scoring functions (Figure 3A right).

For non-recursive and recursive translation models, stochastic decoding methods outperformed deterministic methods on average logP scores (Figure 3A left) and average pairwise diversity (Figure 3B) for generated compounds as a function of recursive iteration. Non-greedy search strategies are not common practice in *de novo* molecular design [Gómez-Bombarelli et al., 2018, Kusner et al., 2017, Jin et al., 2019b]. 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 underemphasized. This

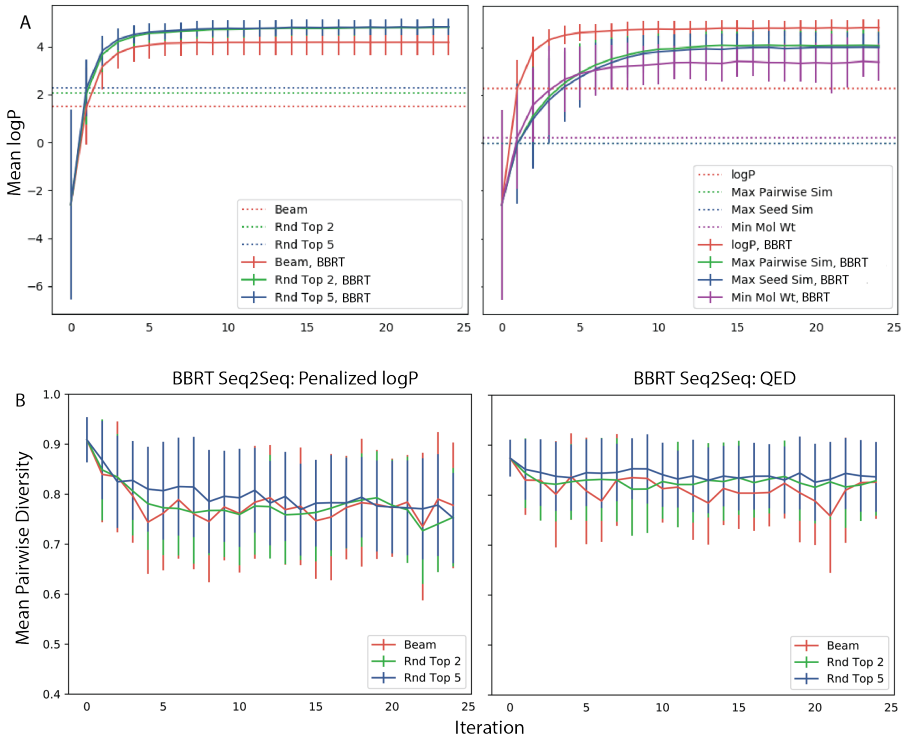

Figure 3: Ablation experiments using BBRT-Seq2Seq. A. (Left): Mean logP from 900 translations as a function of recursive iteration for three decoding strategies. Dotted lines denote non-recursive counterparts. (Right): Mean logP as a function of recursive iteration for four scoring functions. B. (Left): Diversity of generated outputs across recursive iterations for logP translation. (Right): Diversity of generated outputs across recursive iterations for QED translation.

trend has also been observed in the natural language processing (NLP) literature where researchers have recently 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 when using BBRT. When optimizing for logP, we observed that a logP scoring function quickly discovers the best scoring compounds while secondary scoring functions improve logP at a slower rate and do not converge to the same scores. This trade-off highlights the role of conflicting molecular design objectives. For Figure 3A, the standard deviation typically decreased with iteration number n. As property values converge to a certain range, we investigated whether BBRT produces compounds with less diversity. In Figure 3B we show average pairwise diversity of translated outputs per recursive iteration across three decoding strategies and observe decay in diversity for logP. For the best performing decoding strategy, the top-5 sampler, diversity decays from approximately 0.86 after a single translation to approximately 0.78 after $n = 25$ translations. This decay may be a product of the data—higher logP values tend to be less diverse than a random set of compounds. For QED (Figure 3B right), we observe limited decay. Differences in decay rate might be attributed to task variability, one being extrapolative and the other interpolative.

4 Discussion

Motivated by molecular optimization as a translation task, we develop BBRT, a simple algorithm that applies a decoding and scoring strategy at test-time and recursively feeds the output of translation models back into the same model for additional optimization. We apply BBRT to well-known models and produce competitive results with state-of-the-art performance for different property optimization tasks. For future work, we will consider other scoring methods including molecular docking to the target of interest. Lastly, as BBRT is limited by the construction of labeled training pairs, we plan to extend translation models to low-resource settings, where property annotations are expensive to collect.

References

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural Machine Translation by Jointly Learning to Align and Translate. *arXiv.org*, September 2014.

Nicolas Boulanger-Lewandowski, Yoshua Bengio, and Pascal Vincent. Audio Chord Recognition with Recurrent Neural Networks. *ISMIR*, 2013.

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.

Angela Fan, Mike Lewis, and Yann N Dauphin. Hierarchical Neural Story Generation. *ACL*, pages 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.

Alex Graves. Sequence Transduction with Recurrent Neural Networks. *arXiv.org*, November 2012.

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.

Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, November 1997. ISSN 0899-7667.

John J Irwin, Teague Sterling, Michael M Mysinger, Erin S Bolstad, and Ryan G Coleman. ZINC: A Free Tool to Discover Chemistry for Biology. *Journal of Chemical Information and Modeling*, 52(7):1757–1768, 2012.

Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Junction tree variational autoencoder for molecular graph generation. In *International Conference on Machine Learning*, pages 2328–2337, 2018.

Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Multi-resolution Autoregressive Graph-to-Graph Translation for Molecules. *arXiv.org*, June 2019a.

Wengong Jin, Kevin Yang, Regina Barzilay, and Tommi Jaakkola. Learning multimodal graph-to-graph translation for molecule optimization. In *International Conference on Learning Representations*, 2019b.

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.

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*, pages 1945–1954, 2017.

Yujia Li, Oriol Vinyals, Chris Dyer, Razvan Pascanu, and Peter Battaglia. Learning Deep Generative Models of Graphs. *arXiv.org*, March 2018.

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, pages 3104–3112, Cambridge, MA, USA, 2014. MIT Press.

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*, pages 6410–6421, 2018.