Transformer-based Generative Adversarial Networks for Generating Molecules with Desired Properties

Chen Li[1] li260@bio.kyutech.ac.jp
Kazuma Kaitoh [1] kaito168@bio.kyutech.ac.jp
Yoshihiro Yamanishi [1] yamani@bio.kyutech.ac.jp

[1] Kyushu Institute of Technology, 680-4 Kawazu, Iizuka, Fukuoka, 820-8502, Japan

Keywords: Transformer, Generative adversarial network, Policy gradient, Molecular generation

Molecules can be represented by string-based sequences derived from molecular graphs, called the simplified molecular-input line-entry system (SMILES). Generative adversarial networks (GAN) with SMILES strings [1] have attracted widespread attention in generating molecules in drug discovery. Most models apply recurrent neural networks (RNNs) as the generator for the molecular generation with SMILES strings. However, RNNs are difficult to generate molecules with complex rings. In general, highly cyclic molecules have long sequence representations and more strict syntax than acyclic molecules. Slight changes in the syntax may result in the generation of molecules with totally different chemical property, or invalid molecules. Furthermore, RNNs cannot work on GPU versions because the current iteration must compute after the previous time step, which is not conducive to handle big data to explore infinite chemical space.

To overcome the above drawbacks, we propose a transformer-based objective-reinforced GAN model in this study. The model consists of two main parts: *generator* and *discriminator*. The generator is a generative model that tries to generate realistic fake data, and it is a transformer architecture with several stacked encoders and decoders. The discriminator is treated as a binary classifier, attempting to distinguish the generated data to avoid being fooled by the generator. The discriminator is based on a convolutional neural network (CNN), which is composed of a convolutional layer, a max-pooling layer, and a highway layer. Note that the generator and discriminator train in alternation. In addition, a reinforcement learning approach called the Monte Carlo policy gradient (MCPG) [2] is applied. While ensuring that the discriminator effectively guides the training of the generator, it also takes the desired chemical properties into account to generate desired molecules. Concretely, the discriminator first outputs the probability that the current input sequence is from the original SMILES dataset. Then, it calculates the chemical properties of the current input sequence, such as drug-likeness and solubility. Finally, the sum of the probability and the properties are used as a reward for MCPG. In experiments, we test our proposed method on molecular generation from the ZINC chemical dataset, and demonstrate the usefulness of our method in terms of uniqueness, novelty, and diversity in generating molecules.

[1] 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 preprint arXiv:*1705.10843, 2017.
[2] Sutton, R.S.; McAllester, D.A.; Singh, S.P.; Mansour, Y. Policy gradient methods for reinforcement learning with function approximation, *Proc. In Advances in Neural Information Processing Systems 12*, 1057-1063, 1999.