


Augmented Memory: Capitalizing on Experience Replay to Accelerate De Novo Molecular Design

J. Guo^{1,2}, P. Schwaller^{1,2*}

¹Laboratory of Artificial Chemical Intelligence (LIAC), Institut des Sciences et Ingénierie Chimiques, Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland, ²National Centre of Competence in Research (NCCR) Catalysis, Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

Sample efficiency is a fundamental challenge in de novo molecular design. Ideally, molecular generative models should learn to satisfy desired objectives under minimal oracle evaluations (computational prediction or wet-lab experiment). This problem becomes more apparent when using oracles that can provide increased predictive accuracy but impose a significant cost. Molecular generative models have shown remarkable sample efficiency when coupled with reinforcement learning, as demonstrated in the Practical Molecular Optimization (PMO) benchmark. Here, we propose a novel algorithm called Augmented Memory that combines data augmentation with experience replay. We show that scores obtained from oracle calls can be reused to update the model multiple times. We compare Augmented Memory to previously proposed algorithms and show significantly enhanced sample efficiency in an exploitation task and a drug discovery case study requiring both exploration and exploitation. Our method achieves a new state-of-the-art in the PMO benchmark which enforces a computational budget, and outperforms the previous best performing method on 19/23 tasks.



Jeff Guo, Philippe Schwaller, *ChemRxiv*, **2023**. doi: [10.26434/chemrxiv-2023-qmqmq-v3_D_O_I:10.26434/chemrxiv-2023-qmqmq-v3](https://doi.org/10.26434/chemrxiv-2023-qmqmq-v3_D_O_I:10.26434/chemrxiv-2023-qmqmq-v3)