APPENDIX

Anonymous authors
Paper under double-blind review

A ADDITIONAL RESULTS

A.1 SYNTHETIC SCENARIO

As illustrated in Figure 1, the distribution of Top-100 GSK3β scores shows consistent trend in preference-specific GFlowNet and our proposed HN-GFN, although the trend is not significant as the JNK3 property.

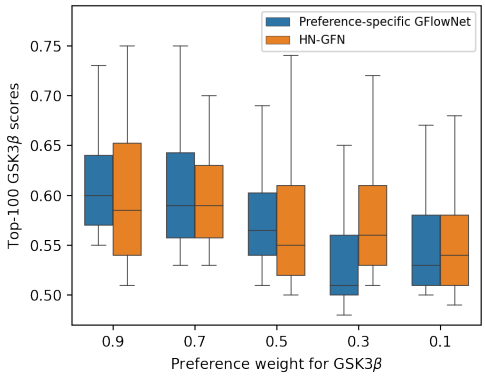

Figure 1: Comparison of the distribution of Top-100 GSK3β scores sampled by different preference vectors using preference-specific GFlowNets and HN-GFN.

A.2 SAMPLED MOLECULES IN MOBO EXPERIMENTS

We give some examples of sampled molecules from the Pareto front by HN-GFN in the GSK3β + JNK3 + QED + SA optimization setting (Figure 2). The numbers below each molecule refer to GSK3β, JNK3, QED, and SA scores respectively.

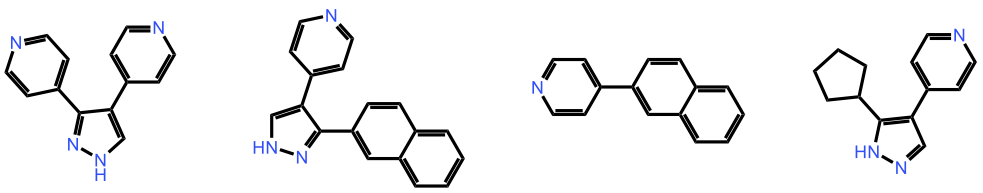

Figure 2: Sampled molecules from the approximate Pareto front by HN-GFN.

B IMPLEMENTATION DETAILS

B.1 MOLECULE DOMAIN

Following (Bengio et al., 2021), the molecules are generated based on a set of 105 building blocks. The same substructure containing multiple stems (atoms for linking another building block) is served as separate building blocks. We allow the GFlowNet generator to sample molecules with 2-8 blocks.

B.2 EXTERNAL TEST SET

First, we generate a random dataset containing 300K molecules uniformly based on the number of building blocks (from 2 to 8). Next, we sample the test sets with uniform property distribution corresponding to GSK3β and JNK3, respectively, from the 300K molecules.

B.3 ORACLE

We adopt the random forest model released by (Jin et al., 2020) as oracles to evaluate the inhibition ability of generated molecules against GSK3β and JNK3.

B.4 BASELINES

All the baselines are implemented using the publicly released source codes with small adaptations for the MOBO scenarios. The GP-BO approach utilizes EHVI as the acquisition function, while HierVAE is implemented in two different settings (qParEGO and qEHVI). For all GP-based methods, each objective is modeled by an independent GP.

B.5 HN-GFN BASED MOBO

The overall MOBO algorithm leveraging HN-GFN as an acquisition function optimizer is described in Algorithm 1. Our proposed algorithm is implemented in PyTorch, and the values of key hyper-parameters are illustrated in Table 1.

Algorithm 1 HN-GFN based MOBO

Input: Oracle O, initial dataset $\mathcal{D}_0 = \{(x_i^0, f(x_i^0))\}_{i=1}^n$, round N, batch size b
for $i = 1$ **to** N **do**
  Fit surrogate model $\mathcal{M}$ on dataset $\mathcal{D}_{i-1}$
  Sample target preference weights Λ
  Train π_θ with HN-GFN
  Sample query batch $B = \{x_1, \ldots, x_b\}$ based on $\lambda_{target} \in \Lambda$
  Evaluate batch B with O: $\hat{D}_i = \{(x_1, O(x_1)\}$
  $\mathcal{D}_{i+1} = \mathcal{D}_i \cup \{(x_1^i, f(x_1^i)), \ldots, (x_b^i, f(x_b^i))\}$
end for

Table 1: Hyper-parameters used in the real-world MOBO experiments.

Hyper-parameter	GSK3β + JNK3	GSK3β + JNK3 + QED + SA
Learning rate (proxy)	2.5e-4	1e-3
Learning rate (GFlowNet)	5e-4	5e-4
Reward exponent	8	8
Trajectories minibatch size	8	8
Offline minibatch size	8	8
Hidden size (molecule)	256	256
Hidden size (preference vector)	100	100

REFERENCES

Yoshua Bengio, Tristan Deleu, Edward J Hu, Salem Lahlou, Mo Tiwari, and Emmanuel Bengio. Gflownet foundations. arXiv preprint arXiv:2111.09266, 2021.

Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Multi-objective molecule generation using interpretable substructures. In International conference on machine learning, pp. 4849–4859. PMLR, 2020.