MMO-GFlowNet: Multi-objective Metahuristic Optimization of GFlowNet Parameters for Molecular Generation

Olaide Nathaniel Oyelade
Computer Systems Technology, College of Science and Technology,
North Carolina A&T University
Greensboro, USA
onoyelade@ncat.edu [0000-0003-1252-0719]

Hui Wang
School of Electronics, Electrical Engineering and Computer Science,
Queen's University Belfast
Belfast, UK
h.wang@qub.ac.uk [0000-0003-2633-6015]

***Abstract*—Generative models have evolved in recent years, and are now capable of learn complex underlying structures of molecules. The structure-based drug design (SBDD) models are a class of generative models that learn the structure of target protein receptors to generate molecules capable of binding to the target. GFlowNets are stochastic-based machine learning models that sample complex objects like protein-protein interactions in drug design, and have recently merged with SBDD. Conditioned GFlowNets are characterized by combinatorial parameters needing carefully selected values, with impaired performance when poorly combined. To address this gap, we propose a novel hybrid of metaheuristics for multi-objective optimization of a subset of conditioned GFlowNet. We specifically adapted a biology-based optimizer to jointly maximize three objectives to learn the combined parameter values suitable to fully train GFlowNet conditioned to pockets of protein structure. Thereafter, drug-like molecules are derived from the model. Using the CrossDocked2020 and docking benchmark version 5 (DB5), we experimentally evaluated the proposed model.**

Keywords—generative model, GFlowNet, drug design, molecular generation, graph, metaheuristic algorithm

I. INTRODUCTION

Structure-based drug design (SBDD) aims to accelerate drug discovery by leveraging three-dimensional protein structures to propose ligands that are chemically valid, synthesizable, and likely to bind a target pocket. In practical pipelines, SBDD is inherently multi-criteria, requiring trade-offs among binding-related proxies (e.g., docking scores), drug-likeness, and synthetic accessibility.

Generative Flow Networks (GFlowNets) provide a principled way to learn a stochastic generative policy that samples objects with probability proportional to a nonnegative reward, enabling diverse goal-directed generation. Recent work has adapted GFlowNets to pocket-conditioned molecular generation for SBDD, including TacoGFN and geometric-informed conditioning variants that incorporate richer pocket–ligand geometric signals [1], [2].

Despite this progress, pocket-conditioned multi-objective generation introduces a practical parameter-selection challenge: reward scaling, objective weighting, and other control parameters materially affect training stability and the trade-off profile of generated candidates. In multi-objective settings, naive scalarization can bias coverage of the Pareto frontier, motivating multi-objective GFlowNet formulations that explicitly target diverse Pareto-optimal solutions [3].

In this paper, we propose MMO-GFlowNet, which treats the selection of key GFlowNet control parameters as a multi-objective optimization problem. MMO-GFlowNet couples a pocket-conditioned GFlowNet with an outer-loop metaheuristic based on the Ebola Optimization Search Algorithm (EOSA) to search for parameter configurations that yield robust multi-objective performance [4], [5].

Our main contributions are as follows:

• A multi-objective meta-optimization framework (MMO-GFlowNet) for selecting GFlowNet control parameters for pocket-conditioned molecular generation.

• An EOSA-driven search procedure that returns Pareto-competitive parameterizations, reducing reliance on ad hoc manual tuning.

• An empirical evaluation on CrossDocked2020 and DB5 demonstrating improved multi-objective trade-offs and sensitivity to key control parameters [6], [7].

II. RELATED WORKS

Pocket-/target-conditioned GFlowNets for SBDD. A key line of work applies GFlowNets to SBDD by conditioning generation on protein pocket structure and optimizing reward-defined objectives. TacoGFN models a reward distribution conditioned on pocket structure and generates molecules with probabilities proportional to affinity- and property-based rewards [1]. Geometric-informed conditioning variants further incorporate richer geometric signals to improve pocket–ligand consistency and downstream docking-related metrics on CrossDocked2020 [8].

Multi-objective generation with GFlowNets. Multi-objective optimization is central to practical molecular design, motivating methods that target coverage of Pareto-optimal trade-offs rather than a single scalarized optimum. Multi-Objective GFlowNets formalize this setting and propose variants designed to generate diverse Pareto-optimal candidates under conflicting objectives, highlighting the importance of candidate diversity alongside Pareto performance [3].

Hybridizing GFlowNets with search/metaheuristics. Several recent methods combine amortized generation with additional stochastic search or heuristic guidance. GFACS

integrates GFlowNets with ant colony optimization to iteratively refine a learned prior distribution via parallel probabilistic search for combinatorial optimization problems [9], [10]. Genetic-guided GFlowNets distill a genetic algorithm into a GFlowNet-trained policy to improve sample efficiency in molecular optimization, with applications including inhibitor design [11]. Related developments in combinatorial domains (e.g., vehicle routing) propose training objectives that combine trajectory-level and local balance signals to improve solution quality and generalization [12].

Protein structure representations. Pocket-conditioned generators rely on informative protein representations. Geometric Vector Perceptrons (GVPs) provide a widely used architecture for learning from protein structures by jointly modelling scalar and vector features on protein graphs, and have been adopted in multiple protein-structure learning pipelines [13].

Our positioning. MMO-GFlowNet is complementary to pocket-conditioned SBDD GFlowNets [1] and to multi-objective GFlowNet formulations [3], but targets a different pain point: systematic multi-objective selection of GFlowNet control parameters. We adopt EOSA as the outer-loop search engine and adapt it to optimize parameter configurations for multi-objective pocket-conditioned molecular generation [4], [5].

III. Methodology

In this section, the proposed approach for multi-objective optimization of GFlowNet parameters is described and illustrated. First, an overview of the proposed method is presented and discussed. Then, the underlying GFlowNet variant and the general compositional process are presented. We proceed to describe the integration of a metaheuristic algorithm into the multi-objective task of finding the best parameter configuration to train the GFlowNet.

A. Overview of MMO-GFlowNet Framework

GFlowNets represent a family of generative machine learning models capable of sampling fragments of compositional systems by optimizing a probability functional proportional to a reward function. Motivated by our recent work [14], in this study, a multi-objective optimization problem is being addressed to improve the quality and diversity of candidate compositional graphs, systems, or objects being generated. The generative tasks considered in this study focus on molecular composition to discover ligands (molecules, drugs) capable of binding to protein receptors to trigger or block a biological process. However, since finding diverse and relevant drugs with high binding affinity is the focus of drug discovery, our proposed approach seeks to pursue this through an optimized configuration of settings for the generative model. In Fig 1, we present an illustration of the entire pipeline of the MMO-GFlowNet, a multi-objective framework that finds optimal settings for GFlowNet.

The figure shows four major components that successfully describe the MMO-GFlowNet method. The first is a Geometric Vector Perceptron with a graph neural network (GVP-GNN) model [13] that provides an embedding representation of the protein receptor (target). The second is a protein receptor target pocket conditioned type of GFlowNet [1], which helps in generating diverse ligands or drugs/molecules that bind to the pockets. Thirdly, the pharmacophore model that learns the 3D composition and chemical properties of molecules. Fourthly, is the metaheuristic algorithm, Ebola optimization search algorithm (EOSA) [4] [5], is applied to find optimal parameter settings for GFlowNet. In this EOSA method is adapted to multiple objectives (three) as opposed to the original single objective method. The docking score, quantitative estimation of drug-likeness (QED), and synthetic accessibility (SA) score are the three objectives jointly optimized by EOSA to discover the optimal parameter configuration for GFlowNet. The following subsections elaborate on the approach described in the overview figure.

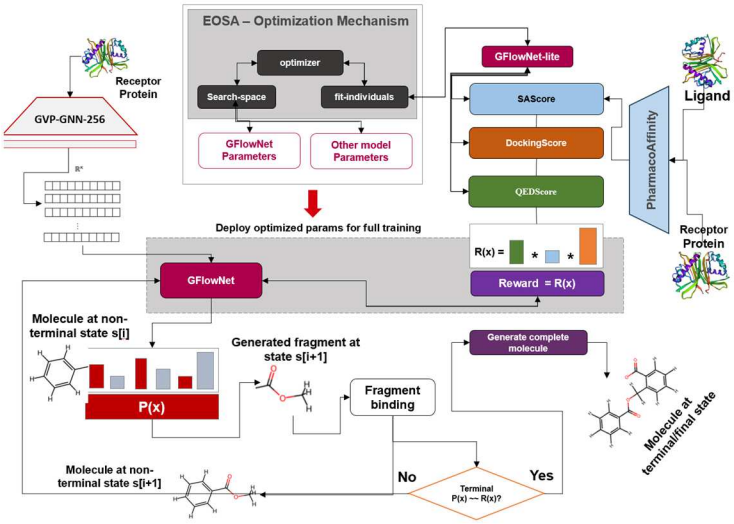

Fig. 1. A framework showing the multi-objective optimization of the parameters of MMO-GFlowNets with the EOSA optimization method

B. Conditioned GFlowNet

Most applications of GFlowNet are typically different from the general GFlowNet, which is domain-agnostic, whereas the conditional GFlowNet is adapted to learn a family of reward functions. In this study, protein receptor pocket structure conditioned GflowNet is the baseline of the generative method used for the multi-objective method described. We particularly extend our recent work [14] and TacoGFlowNet [1]. General GFlowNets aims to find diverse representations of potential transitions leading to candidate compositional objects through an acyclic graph $G = (S, \tau)$. The graph G, transition over several states $S = \{s_o, \ldots . s_n\}$ through some actions $Å = \{a_o, \ldots . a_{n-1}\}$ that aims to optimize the proportionality of the probability $\pi(x)$ outcome with the reward function $R(x)$. As a result, τ is a transition which draws from $\{s_0 \rightarrow s_1 \ldots, s_t \rightarrow s_{t+1}, \ldots s_{n-1} \rightarrow s_n\}$. This kind operation produces several trajectories Trooted in s_0. With T_i representative of a potential molecular graph. Each $\tau(t)$ requires an action a_t, to produce s_t, preferably, denoted by s to produce s', and jointly defined in equation (1).

$$s' = \{\tau(s, a) \mid s \rightarrow s' \in \tau \; and \; a \in Å\} \qquad (1)$$

The conditioning of the GFlowNet described in this study seeks to optimize a stochastic or probability function $\pi(x|Pckt)$

that is conditioned to the embedding of the protein receptor pocket to be proportional to the reward function $R(x|Pckt)$. Moreover, the pocket conditioning also includes the exponent β to the reward function, which seeks to direct the generative model towards the reward function. Therefore, the proportionality between the stoachsitci function and the reward function will follow $\pi(x|Pckt,\beta) \propto R(x|Pckt,\beta)$. The targeted outcome producing diverse samples of $\{s_0 \rightarrow s_1 \ldots, s_t \rightarrow s_{t+1}, \ldots s_{n-1} \rightarrow s_n\}$ that produces a drug-like molecule/ligand capable of strongly binding with protein receptor pocket $Pckt$ so that any s_t is a potential fragment of molecular composition required for the full molecular graph. However, with the multi-objective metaheuristic optimization propose in this study, the proptionality is slightly adjusted to $\pi(x|Pckt,\beta,\theta^*) \propto R(x|Pckt,\theta^*)^{\beta}$. Considering the sharp closeness of this proportionality of the stochastic function with the reward function in [1], we also leveraged their drug-like molecular fragment vocabulary. Their vocabulary was built using the BRICS (Breaking of Retrosynthetically Interesting Chemical Substructures) decomposition on curated subset of 250,000 small, drug-like molecules extracted from the ZINC database. Additionally, the probability π is derived using their approach of scoring the likelihood of s_t similar to $Pckt$ at every $\tau(t)$ by using a graph transformer.

In Fig 2, we illustrate the use of the π at transition $\tau(t)$ in the pathway of an arbitrary trajectory T_t . Therefore, taking trajectory $\{s_0 \rightarrow s_1 \ldots, s_t \rightarrow s_{t+1}, \ldots s_{n-1} \rightarrow s_n\}$ at location i, of $\pi(T_i)$ will be proportional to the reward $R(T_i)$.

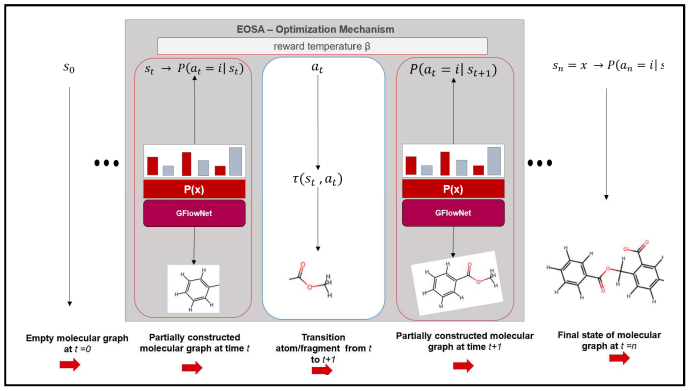

Fig. 2. An illustration of the computation of the likelihood function π at an arbitrary time (t) showing transition $\tau(t)$

With the transition function previously defined, the reward function $R(T_i)$ is derived from three separate metrics that support the conditioning of the generative model. These metrics are the quantitative evaluation of drug-likeness (QED score (g), synthetic accessibility (v), and affinity docking score (ℓ). Values returned from g, v, and ℓ are drawn from the ranges [0, 1], [1, 10], and [-4, -12]. These three jointly evaluate the affinity property, drug-likeness property, and synthesizability property of the molecular graph generated, and are derived in equation (2):

$$R(x) = R(T) = g * v * \ell \tag{2}$$

At the core of the training process is learning a learning a loss function that reflects the trajectory balance objective function. This function combines the forward and backward pass over some batches of input. The forward pass computes a probability value $P_f(\acute{x}|Pckt,\beta,\theta^*)$ for some partially constructed molecular graph $\acute{x}$ at state s'. Similarly, the backward pass computes value $P_b(\acute{x}|Pckt,\beta,\theta^*)$ for the same $\acute{x}$ back to state s. As a result, the trajectory function considered during the training of the MMO-GFlowNet model is described in equation (3):

$$L(\tau) = \left(\log \frac{Z(Pckt,\beta,\theta^*) * \prod_\tau P_f(\acute{x}|Pckt,\beta,\theta^*)}{R(x|Pckt,\beta,\theta^*) * \prod_\tau P_b(\acute{x}|Pckt,\beta,\theta^*)}\right)^2 \tag{3}$$

Where Z is a partition function, is also a sum of all the rewards R, and θ^* is the optimized learnable parameters. The outcome of this learning process is a stochastic policy that will be proportional or approximately the ratio of the reward and partition function: $\pi \approx \frac{R(x|Pckt)^{\beta}}{Z(Pckt,\beta)}$.

In the next subsection, we describe the initial optimization of a subset of those parameters found in the learnable parameter set θ^*.

C. *Multi-objective Parameter GFlowNet optimization*

Metaheuristic algorithms have been recently integrated with GFlowNet for trajectory balance in both the forward and backward directions. In his study, we approach the use of the metaheuristic algorithm to support GFlowNets in finding the optimal parameter settings to kick-start the training process. The aim is to reduce the struggle and challenges resulting from poorly selected parameters. Considering also that the use of these parameters in the learning process of GFlowNets often requires a combinatorial selection, we creatively adapted the EOSA algorithm to discover optimal configuration settings. The proposed use of EOSA for jointly optimizing the values of the docking score, QED score, and the synthesizability access score represents the multi-objective tasks described in the scope of this study. This approach differs from the inherent GFlowNet optimization introduced in Multi-Objective GFlowNet (MO-GFN) [3] whereas their work is aimed at conditioning the GFlowNet on multi-objective preference, this study applies the multi-objective optimization outside the GFlowNet model and to the initial parameter configuration.

Given a set of parameter configurations, Θ for training GFlowNets, we seek to train a metaheuristic algorithm to find an optimal $\theta \in \Theta$ by maximizing multiple objectives that unanimously enhances the performance of the GFlowNets model to produce diverse candidates as illustrated in equation (4):

$$\max_{\theta \in \Theta} M(\theta, [\ell, g, v]) \tag{4}$$

Note that the ℓ, g, and v are notations previously defined in the last subsection. Contrary to the multi-objective optimization described in [3], which aims to directly maximize conflicting

objectives within the GflowNet model, our proposal takes the task of optimization to another surrogate optimizer, namely the EOSA metaheuristic algorithm. As a result, the best and optimal configuration of parameters θ, is then applied to fully train the GFlowNet model.

The EOSA algorithm was inspired from the propagation strategy of the Ebola virus. The algorithm leverages of the natural compartments of sub-populations: Update of Susceptible (S), Infected (I), Hospitalized (H), Exposed (E), Vaccinated (V), Recovered (R), and Quarantine (Q). In utilizing the EOSA method, a set of parameters was identified from the GFlowNet configuration and was collected into $\chi \in \Theta$. We identified leaky coefficient, reward multiplier, highest possible QED reward, highest possible docking score reward, highest possible synthetic accessibility reward score, QED exponent, synthetic accessibility exponent, global batch size, and learning rate decay as the list of parameters for the combinatorial search problem. Therefore, each of these parameters describes into $\chi \in \Theta$, and every χ has an upper and lower value, seen in equation (5), which bounds the EOSA search process.

$$\mathrm{P} = \{\, \chi[\mathrm{i}]_b^{\upsilon} \mid \upsilon, b \in \mathbb{R} \wedge \mathbb{Z}, i \in 0 \ldots N\} \tag{5}$$

The population, that is the susceptible compartment, is initialized using equation (6), which derives from a permutation of that the $\chi[0]^{\upsilon}$ or $\chi[0]_b \ldots . \chi[\mathrm{N}-1]^{\upsilon}$ or $\chi[\mathrm{N}-1]_b$:

$$\mathrm{X} = \frac{|P|!}{(|P|-2)!} \tag{6}$$

The problem size for the EOSA optimization is dynamic and computed from $|X|$. Moreover, EOSA computes the fitness value of each element in X using a weighted sum of the three objective scores as described in equation (7) and φ denotes a lite evaluation on GFlowNet:

$$\mathcal{F} = \varphi(\omega_1 \ell + \omega_2 \not{g} + \omega_3 \nu) \tag{7}$$

Note that the ω_1, ω_2, and ω_3 are assigned values according to their contribution to the whole relevance of elements in the population. During experimentation, we assigned the values of each of these weights. Since the aim of the optimization is to find the best combination of parameter values suitable for fully trainingGFlowNet, during each iteration, the best solution is derived based on the fitness function. Subsequent iterations only need to update the global best (gb) by comparing the fitness value ($\mathcal{F}$) of the current best (cb) with the fitness ($\acute{\mathcal{F}}$) of the gb in the search space. This is described in equation (8):

$$gb = \begin{cases} gb, & \acute{\mathcal{F}} < \mathcal{F} \\ cb, & \acute{\mathcal{F}} \geq \mathcal{F} \end{cases} \tag{8}$$

In Algorithm 1, we demonstrate the pseudocode of the multi-objective optimization using the EOSA algorithm method. The algorithm requires three inputs, namely the parameters with upper and lower bounds, which are combined to form the search space (Θ), the duration (t) for training GFlowNet to obtain the fitness value, and the number of iterations (*iter*). The output of the algorithm is an optimized subset of the algorithm, replaceable in the complete set of parameters for full training of the GFlowNet.

Algorithm 1: MMO-GFlowNet Algorithm

Data: $\Theta, t, iter$; /* params to optimize, fit-time, & iter */
Result: Θ^{*}_{opt}; /* optimized params */
$X \leftarrow []$;
for $i = 0$ *to* $size(\Theta) - 1$ **do**
  $u, l \leftarrow \Theta[i]_0, \Theta[i]_1$;
  $X[i] \leftarrow (u, l)$;
end
$S \leftarrow {}_nP_r(X)$; /* search space */
$F \leftarrow GFlowNet(S, t)$; /* fit search space */
$Q \leftarrow [S, F]$;
for $i = 1$ *to* $iter$ **do**
  $Q' \leftarrow EOSA(Q)$;
  $Q \leftarrow Q'$;
end
$S, \_ \leftarrow sort(Q, \nabla)$;
$\Theta^{*}_{opt} \leftarrow S[0]$;
return Θ^{*}_{opt};

This section summarizes the GFlowNet conditioned on the pocket structure of the protein receptor, and also the multi-objective optimization of parameters for fully training it. In the following sections, we describe the experimentation and results

IV. Experimentation and Dataset

The experimental setup, which describes the computational environment, is presented alongside the datasets applied to the proposed method.

A. Parameter settings and system setup for experimentation

The assignment of values to all the control parameters in the algorithm for the proposed method is described in Table I. Outlining these values allows for reproducibility of the proposed method. The assignment of the upper and lower values was inspired by our previous experiment reported in [14]. The training time of 50 epochs was chosen for the training of EOSA and GFlowNet to minimize the complexity of the optimization process. In addition, the values assigned to the EOSA parameters were also listed in the table.

TABLE I. Parameter definition and value assignments

Param	**MO-GFlowNet**		
	Description	***Upper value***	***Lower value***
LC	Leaky coefficient	0.1	0.3
RM	Reward multiplier	2.0	4.0
MQR	Maximum Quantitative estimation of drug-likeness reward	0.6	0.8
MDR	Maximum value for the _docking _reward	-4.0	-0.6
MSR	Maximum synthetic accessibility reward	0.7	0.9
QED-E	Quantitative estimation of drug-likeness exponent	0.85	1.0
SA-E	Synthetic accessibility exponent	0.60	1.0
GBS	Global batch_size	2	16
LR-D	Learning rate decay	19000	30000
	Optimization training time		

Param	MO-GFlowNet				
	Description			*Upper value*	*Lower value*
Epoch	EOSA-time	50	**Epoch**	GFlow Net-Lite	50
	EOSA parameters				
Param	**Description**			**Values**	
S	Problem size(susceptible population size)			$\|X\|$	
ε	Epsilon value			0.0001	
β and π	Contact and recruitment rate			0.1 and 0.1	

All training was carried out in a computational environment with Nvidia RTX 6000 Ada Generation, 48 GB GDDR6, 64GB, memory, Intel Xeon w7-3565X (82.5MB Cache, 32 cores, 64 threads, 2.5GHz to 4.8GHz Turbo 335W), and 2TB Performance SSD, SED Ready, units based on the A100 Tensor Core GPU with memory size of 80 GB. Implementation was done with Pytoch using the Python 3 programming language alongside other dependencies.

B. *Datasets*

Two publicly available datasets were applied to the experimentation. The CrossDocked2020 [6] dataset and docking benchmark version 5 (DB5) [7] datasets. The CrossDocked2020 contains 52,126,979 files and is a large-scale, open-source dataset designed for training machine learning models in structure-based drug design (SBDD), specifically for predicting protein-ligand binding poses and affinities. It is widely used to train AI models to recognize ligands in target structures. On the other hand, the DB5 has three classifications of cases, namely the rigid, medium, and difficult. This classification demonstrates the level of conformation to changes when binding takes place. The rigid category has 162 samples, the medium category has 60, and the category named difficult has 35 samples, thereby totaling 257. These two datasets were separately applied to the proposed method, and result obtained are described in the next section.

V. Results and Discussion of Findings

In this section, results are presented and evaluated comparatively for the optimization task and generative model. Findings from the results are discussed in the context of generating molecules using the structure-based drug design approach.

Table II outlines the results obtained for three of the parameters we optimized using the proposed method on the CrossDock2020 dataset. Whereas the baseline GFlowNet in [8] utilizes the values of 0.7, -0.5, and 0.8 for MQR, MDR, and MSR, respectively, our proposed method discovered the optimal combined values for these parameters to be 0.8, -0.4, and 0.9, respectively. This global best solution was selected through the optimization training, which returned the fitness value of 1.67E+0 for the global best combination of the parameter values. In Fig 3, we show the search trajectory in the search space and the corresponding population state for all the sub-groups/compartments (S, I, V, H, Q, and D) in the EOSA method. The search trajectory reveals an interesting pattern in the search space.

TABLE II. Result showing the fitness and optimized parameters comparable with the default GFlownet model

Dataset	Method	MQR	MDR	MSR	Fitness
Crossdock2020	GFlowNet	0.7	-0.5	0.8	N/A
Crossdock2020	MMO-GFlowNet	0.8	-0.4	0.9	1.67E+0

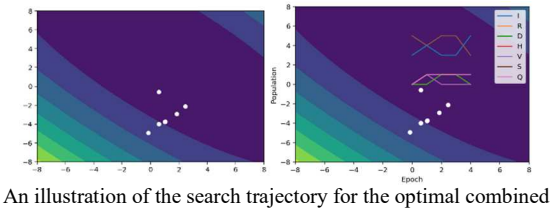

Fig. 3. An illustration of the search trajectory for the optimal combined parameters for MMO-GFlowNet, and the distribution of values to the sub-population groups in EOSA

To evaluate the quality of the solutions in the search space, after the complete training of the EOSA method on the parameter states of GFlowNet, Table III provides the values obtained. Specifically, values for fitness revealing the best, mean/average, standard deviation (std), worst (wst), median, and deviation (dev) were collected among all the potential candidates in the search space. Results for these variables are outline on the table with the best fitness derived as 1.67, and a deviation of -1.67.

TABLE III. Result showing the performance of the EOSA optimizer in finding the best, mean, STD, worst (WST), median (MED) and deviation (DEV)

Model	Dataset	best	mean	std	Wst	med	Dev
MMO-Gflow Net	Crossdock2020	1.67	1.67	2.22 E-16	1.67	1.67	-1.67

The best solution obtained from the optimization process was deployed to train the MMO-GFlowNet method, and the results obtained are listed in Table IV. We evaluated both baseline GFlowNet and MMO-GFlowNet using the prediction time, prediction, QEDs, synthetic accessibility score (SAS), diversity, and novelty metrics. In most of the metrics evaluated, except for prediction time and prediction, the MMO-GFlowNet demonstrated good performance even with the partial training applied to it. This confirms the positive influence of the multi-objective optimization method.

TABLE IV. Result showing mean and median values for the pre-optimized and optimized states of the GFlowNets based on the metrics evaluating the molecular-generated on Crossdock2020

	GFlowNet		MMO-GFlowNet	
Metrics	**Mean**	**Median**	**Mean**	**Median**
Time	5.21281	5.24176	6.43812	6.55227
Prediction	6.14702	6.27868	-6.14735	-6.32176
QEDs	0.383947	0.386194	**0.386794**	**0.389277**

	GFlowNet		MMO-GFlowNet	
Metrics	**Mean**	**Median**	**Mean**	**Median**
SAS	0.656636	0.656983	**0.657458**	**0.658061**
Diversity	0.855935	0.858629	**0.857944**	**0.860829**
Novelty	0.544535	0.546813	**0.548161**	**0.550661**

Figures 4 and 5 show the plot of the prediction values against the novelty of molecules generated, the drug-likeness (QED) of the molecules, and the likelihood of synthesizing the molecules in a wet lab. The emphasis here is on the variation in the plots for the three metrics (Novelty, SAS, and QED) compared between baseline GFlowNet and MMO-GFlowNet. In all three plots for novelty, SAS and QED, MMO-GFlowNet demonstrates better convergence than baseline GFlowNet. An indication of consistency in its performance across all data instances represented on each plot.

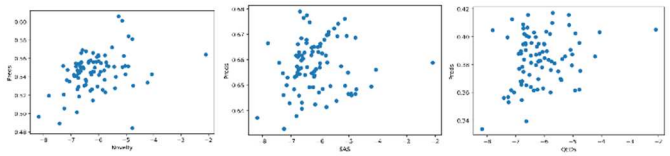

Fig. 4. Scatter chart showing the prediction against novelty, SAS, and QED based on the GFlowNet model on the CrossDock202 dataset

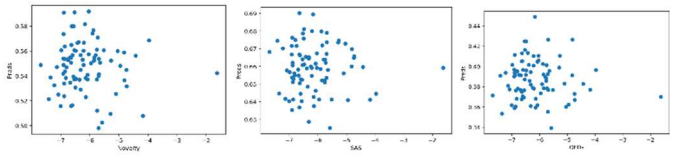

Fig. 5. Scatter chart showing the prediction against novelty, SAS, and QED based on the GFlowNet model on the CrossDock202 dataset

Furthermore, both the GFlowNet and MMO-GFlowNet were applied to generate molecules. Figures 6 and 7 show the outcome of the molecules generated by the two separate methods. We observed that the QED values for the MMO-GFlowNet were higher than those of the baseline GFlowNet. Moreover, complex molecular graphs generated by the MMO-GFlowNet returned higher values of QED compared with the baseline GFlowNet. The implication of this is that MMO-GFlowNet is capable of generating both small and complex molecular graphs at a desired drug-likeness – a feature required in drug design.

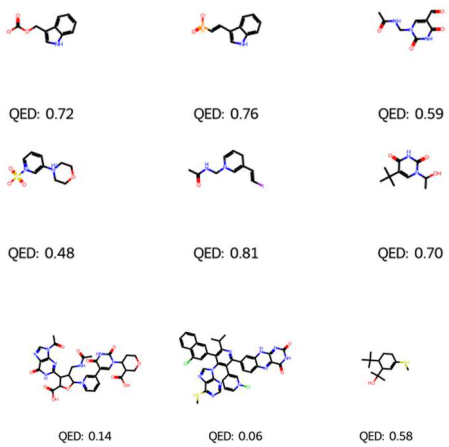

Fig. 6. Sample molecules generated with corresponding QED values using the GFlowNet model on the CrossDock202 dataset

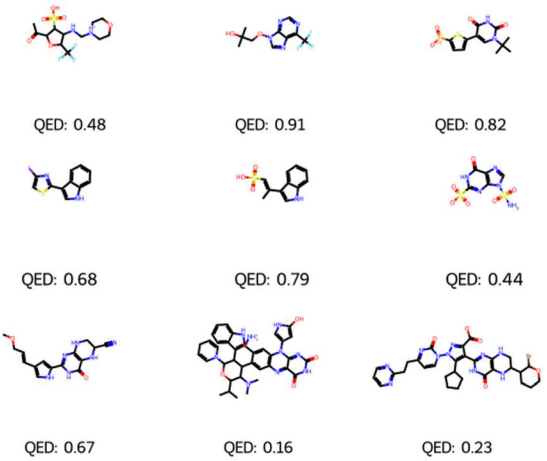

Fig. 7. Sample molecules generated with corresponding QED values using the MMO-GFlowNet model on the CrossDock202 dataset

Findings from the results show the suitability of the MMO-GFlowNet over simply fully training GFlowNets with guessed parameters. It confirms that combinatorially evaluating parameter values before using them for training the generative model might help produce a stronger and more useful generative model applicable to support the SBDD approach. This underlying benefit of this is the richness in diversity of the molecules being generated, and the attractiveness of the novelty and drug-likeness of these generated molecules using methods such as the proposed. MMO-GFlowNet.

VI. Conclusion

This study presents a new approach to the use of a traditional metaheuristic optimization algorithm in the recent GFlowNet method. The hybridization of the two methods is aimed at finding an optimal combination subset of GFlowNet parameters suitable for fully training it for better performance. As a result, the EOSA optimization algorithm was applied to a multi-objective task for addressing a combinatorial parameter selection problem. The optimized parameters were deployed to fully train the GFlowNet model for generating diverse molecular graphs and achieving higher rewards. We experimentally tested the proposed method on the CrossDock202 and DB5 datasets. Results show that the proposed hybrid method contributed to the performance of the GFlowNet model, thereby improving the quality of the molecules generated. Evaluation was done using the docking score, QED score, and synthesizability access score. In the future, we recommend investigating the role of the pharmacophore model to validate the scoring function. Also, it might be relevant to investigate other transformer-based embedding models to represent pocket structures.

VII. References

[1] T. Shen, S. Seob, G. Leea, M. Pandeye, J. Smithe, A. Cherkasove, W. Y. Kimb and M. Ester, "Target-conditioned GFlowNet for Structure-based Drug Design," *Transactions on Machine Learning Research,* 2024.

[2] G. Lee, T. Shen and M. Ester, "Geometric-informed GFlowNets for Structure-Based Drug Design," in *MoML 2024* , 2024.

[3] J. Moksh, R. S. Chandra, H.-G. Alex, R.-B. Jarrid, B. Yoshua, M. Santiago and B. Emmanuel, "Multi-objective gflownets," in *International Conference on Machine Learning*, 2023.

[4] O. N. OYELADE, A. E.-S. EZUGWU, T. I. A. MOHAMED and L. ABUALIGAH, "Ebola Optimization Search Algorithm: A New Nature-Inspired Metaheuristic Optimization Algorithm," *IEEE Access,* vol. 10, pp. 16150-16177, 2022.

[5] O. N. Ovelade and A. E. Ezugwu, "Ebola Optimization Search Algorithm: A new nature-inspired metaheuristic algorithm for global optimization problems," in *2021 international conference on electrical, computer and energy technologies (ICECET)*, 2021.

[6] P. G. Francoeur, T. Masuda, J. Sunseri, A. Jia, R. B. Iovanisci, I. Snyder and D. R. Koes, "Three-Dimensional Convolutional Neural Networks and a Cross-Docked Data Set for Structure-Based Drug Design," *Journal of Chemical Information and Modeling,* vol. 60, no. 9, 2020.

[7] T. Vreven, I. Moal, A. Vangone, B. Pierce, P. Kastritis, M. Torchala, R. Chaleil, B. Jimenez-Garcia, P. Bates, J. Fernandez-Recio, A. Bonvin and Z. Weng, "Docking benchmark version 5 and affinity benchmark version 2.," *Journal of Molecular Biology,* vol. 427, pp. 3031-41, 2015.

[8] T. Shen, S. Seo, G. Lee, M. Pandey, J. R. Smith, A. Cherkasov, W. Y. Kim and M. Ester, "TacoGFN: Target-conditioned GFlowNet for Structure-based Drug Design," *Transactions on Machine Learning Research,* 2024.

[9] M. Kim, S. Choi and H. Kim, "Ant Colony Sampling with GFlowNets for Combinatorial Optimization," *arXiv:2403.07041v2 [cs.LG],* 2024.

[10] M. Kim, S. Choi, H. Kim, J. Son, J. Park and Y. Bengio, "Ant Colony Sampling with GFlowNets for Combinatorial Optimization," in *Proceedings of Machine Learning Research - 28th International Conference on Artificial Intelligence and Statistics, AISTATS 2025*, 2025.

[11] H. Kim, M. Kim, S. Choi and J. Park, "Genetic-guided GFlowNets for Sample Efficient Molecular Optimization," in *NIPS '24: Proceedings of the 38th International Conference on Neural Information Processing Systems*, 2024.

[12] N. Zhang and Z. Cao, "Hybrid-Balance GFlowNet for Solving Vehicle Routing Problems," in *NeurIPS 2025*, 2025.

[13] B. Jing, S. Eismann, P. Suriana, R. J. Townshend and R. O. Dror, "LEARNING FROM PROTEIN STRUCTURE WITH GEOMETRIC VECTOR PERCEPTRONS," in *ICLR 2021*, 2021.

[14] O. N. Oyelade and H. Wang, "A GFlowNet-Based World Model for Structural Protein–Protein Interaction Prediction on DB5," in *Neural Information Processing, ICONIP 2025*, 2025.

[15] B. Li, Z. Wang, Z. Liu, Y. Tao, C. Sha, M. He and X. Li, "DrugMetric: quantitative drug-likeness scoring based on chemical space distance," *Briefings in Bioinformatics,* vol. 25, no. 4, 2024.

[16] P. Ertl and A. Schuffenhauer, "Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions," *Journal of Cheminformatic,* vol. 1, 2009.

[17] O.-E. Ganea, X. Huang, C. Bunne, Y. Bian, R. Barzilay, T. Jaakkola and A. Krause, "Github - equidock_public," [Online]. Available: https://github.com/octavian-ganea/equidock_public/tree/main/data/benchmark5.5/cv/cv_0. [Accessed 10 5 2025].

[18] X. Ji, X. Zhang, W. Xi, H. Wang, O. Gadyatskaya and Y. Li, "Meta Generative Flow Networks with Personalization for Task-Specific Adaptation," *Information Sciences,* vol. 672, 2024.