Synergy of GFlowNet and Protein Language Model Makes a Diverse Antibody Designer

Mingze Yin[1,2*], Hanjing Zhou[1*], Yiheng Zhu[1*], Jialu Wu[3], Wei Wu[4], Mingyang Li[5], Kun Fu[5], Zheng Wang[5], Chang-Yu Hsieh[3†], Tingjun Hou[3†], Jian Wu[1,2,6†]

[1]College of Computer Science & Technology and Liangzhu Laboratory, Zhejiang University
[2]State Key Laboratory of Transvascular Implantation Devices of The Second Affiliated Hospital, Zhejiang University
[3]College of Pharmaceutical Sciences, Zhejiang University
[4]School of Artificial Intelligence and Data Science, University of Science and Technology of China
[5]Alibaba Cloud Computing
[6]Zhejiang Key Laboratory of Medical Imaging Artificial Intelligence
{mzyin256, zhj85393}@gmail.com, {zhuyiheng2020, jialuwu, kimhsieh, tingjunhou, wujian2000}@zju.edu.cn
{sangheng.lmy, fukun.fu, wz388779}@alibaba-inc.com, urara@mail.ustc.edu.cn

Abstract

Antibodies defend our health by binding to antigens with high specificity, primarily relying on the Complementarity-Determining Region (CDR). Yet, current experimental methods of discovering new antibody CDRs are heavily time-consuming. Computational design could alleviate this burden, with protein language models demonstrating remarkable utility in many recent studies. However, most existing models solely focus on antibody developability and struggle to encapsulate the diverse range of plausible CDR candidates, limiting their effectiveness in real-world scenarios as binding is only one factor in the multitude of drug-forming criteria. In this paper, we introduce PG-AbD, a framework uniting Generative Flow Networks (GFlowNets) and pretrained Protein Language Models (PLMs) to successfully generate highly potent, diverse and novel antibody candidates. We innovatively construct a Products of Experts (PoE) composed by the global-distribution-modeling PLM and the local-distribution-modeling Potts Model to serve as the reward function of GFlowNet. The joint training paradigm is introduced, where PoE is trained by contrastive divergence with the negative samples generated by GFlowNet, and then guides GFlowNet to sample diverse antibody candidates. We evaluate PG-AbD on extensive antibody design benchmarks. It significantly outperforms existing methods in diversity (13.5% on RabDab, 31.1% on SabDab) while maintaining optimal developability and novelty. Generated antibodies are also found to form stable and regular 3D structures with their corresponding antigens, demonstrating the great potential of PG-AbD to accelerate real-world antibody discovery.

1 Introduction

Antibodies, also known as immunoglobulins, are specialized proteins produced by the immune system to safeguard and combat a variety of diseases by specifically binding to antigens (Zhu et al. 2022). As shown in Figure 1, an antibody is a Y-shaped molecule composed of two identical heavy chains and two identical light chains linked by disulfide bonds. On each chain, the Complementarity-Determining Region (CDR), also known as the hypervariable region, primarily determines the antibody's affinity and specificity for its target. Given the crucial role of CDRs in antigen binding, computational antibody design primarily focuses on designing these regions. However, this presents a significant challenge due to the enormous combinatorial search space of over $\mathcal{O}(20^{60})$ possible CDRs and the limited solution space satisfying the desired constraints of binding affinity, stability, and synthesizability (Raybould et al. 2019).

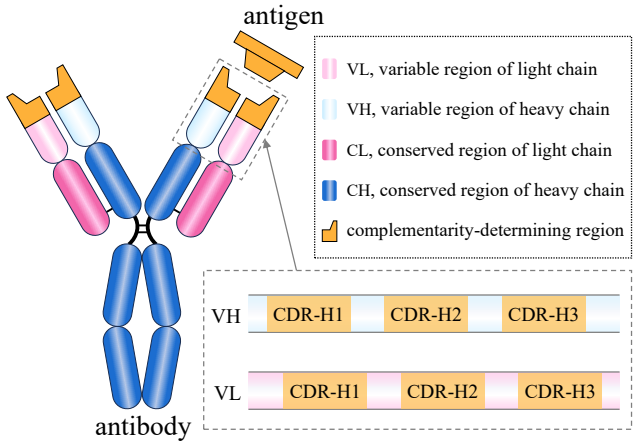

Figure 1: Overview of antibody Y-shape structure.

Current antibody discovery procedures still heavily rely on a laborious combination of high-throughput screening and experimental heuristics. The impressive progress of deep learning methods has recently enabled computational predictions (Jin et al. 2021; Fu and Sun 2022) that could potentially improve wet-lab antibody discovery. Despite the versatility and claimed high prediction accuracy (often defined as "reconstructing the ground truth in training or test datasets") of contemporary deep learning models, there remains room for improvement. Rather than solely focusing

*These authors contributed equally.
†Corresponding authors.

on prediction accuracy, we advocate considering three crucial factors simultaneously when evaluating generated antibody candidates: (1) *developability*: This is a key factor for efficient and economical antibody discovery. Evaluating a candidate's ability to be developed into a real, functional antibody reduces wasted resources in downstream experimental validation. (2) *diversity*: Diversity is essential in real-world design due to the variety of biological targets, binding modes, and structures. Moreover, in antibody design, sequence similarity does not always guarantee similar binding behavior. During ideation, the goal is to produce antibody candidates as diverse as possible, facilitating multi-attribute evaluation to identify antibodies with distinct functionalities. Therefore, generating diverse CDR candidates is a significant and challenging aspect for computational antibody design. (3) *novelty*: Ideally, antibody discovery aims to identify not only effective binders but also those with novel antigen recognition mechanisms. This is crucial for developing vaccines against new or resistant pathogens. While achieving "out-of-distribution generation" is challenging, models should strive to create novel antibody candidates with unique binding modes outside the training data set.

To better incorporate the aforementioned factors in a generative model, exploring new techniques is worthwhile. Protein Language Models (PLMs) have been significantly successful, achieving state-of-the-art performance in various protein-related tasks, including structure prediction (Rives et al. 2021; Lin et al. 2023), evolution-aware prediction (Wang, Ye, and Zhou 2023), and protein function prediction (Xu et al. 2023). However, in this paper, we argue that prevailing approaches, directly applying pretrained PLMs for antibody design, are not at the their best in generating diverse and novel antibody CDRs because low-temperature sampling obstructs the ability to capture the extensive variety of plausible solutions. This observation leads to an important yet challenging research question: *"Can we leverage powerful pretrained PLMs to generate diverse and novel CDR candidates for antibody discovery?"*

In this paper, we propose the **PG-AbD** framework, which unifies **P**rotein language model and **G**enerative flow network for computational **A**nti**b**ody **D**esign. We leverage the Generative Flow Networks (GFlowNets) to fully unleash the potential of pretrained PLMs, adapting them for antibody CDR design while simultaneously ensuring the developability, novelty, and diversity. The pretrained PLM evaluates antibody sequences sampled by GFlowNet and guides its subsequent exploration. Alternatively, GFlowNet thoroughly explores the target antibody distribution space, efficiently generating a diverse set of antibody CDR sequences with high rationality as indicated by their reward scores.

Our PG-AbD framework avoids training data dependency thanks to GFlowNet's exploration capabilities. This allows it to tackle both novel antibody design and antibody optimization. For novel design: GFlowNet can independently explore and identify rational design trajectories for antibodies, even without additional training data. This opens the door for using PLMs to design completely new antibodies with desired functions, even when limited antibody data is available for LLM pretraining and fine-tuning. Additionally, to manage the vast search space (over $\mathcal{O}(20^{60})$ possible CDRs), the template CDR with the desired function is incorporated into the policy network using the FiLM (Vincent et al. 2018) feature fusion module. For antibody optimization: When sufficient training sequences are available, GFlowNet leverages both exploration and exploitation. We combine a fixed PLM with a trainable Potts Model (Hopf et al. 2017; Russ et al. 2020) to construct a Products of Experts (PoE) (Hinton 1999, 2002) reward function for GFlowNet. This PoE integrates global constraints from general protein domains with local constraints specific to antibody domains. GFlowNet and PoE are trained jointly, with PoE exploiting contrastive divergence and negative samples provided by GFlowNet through an MCMC transition proposal. Meanwhile, GFlowNet employs trajectory balancing, utilizing training data trajectories and exploring optimized antibody trajectories proportional to the PoE probability.

In summary, our main contributions are as follows:

- We introduce PG-AbD framework, effectively leveraging the synergy of GFlowNet and PLM to construct a diverse antibody designer.
- PG-AbD is training-data-free. GFlowNet can operate in a self-exploring mode guided by the pretrained PLM, even without training data. Alternatively, GFlowNet can explore and exploit based on existing data trajectories, enabling joint training with a Products of Experts constructed from the pretrained PLM and a Potts Model.
- PG-AbD demonstrates promising performance on multiple antibody design benchmarks. PG-AbD achieves optimal developability and novelty while significantly improving the diversity of generated candidates (13.5% on RabDab and 31.1% on SabDab), offering a valuable tool for future real-world antibody discovery endeavors.
- We execute the complex structural binding simulation of generated antibodies and corresponding antigens, obtaining regular and stable antigen-antibody 3D structures.

2 Related Work

2.1 Antibody design

The design of antibodies primarily focuses on generating CDRs, which are the most important regions determining the binding affinity. Traditional computational methods (Li, Pantazes, and Maranas 2014; Adolf-Bryfogle et al. 2018) iteratively optimize the energy of antibodies to achieve the most stable states. Deep generative models (Shin et al. 2021; Fu and Sun 2022; Luo et al. 2022) have also been adopted as the antibody designer. And multiple GNN-based methods are proposed to co-design antibody sequences and structures (Jin et al. 2021; Jin, Barzilay, and Jaakkola 2022; Kong, Huang, and Liu 2023). Pioneering works (Melnyk et al. 2023; Gao et al. 2023) have explored to employ powerful LLMs for antibody design, but fail to simultaneously balance the developability, novelty, and diversity.

2.2 GFlowNets (Generative Flow Networks)

Initially inspired by the reinforcement learning, (Bengio et al. 2021) lays out the theoretical foundations of

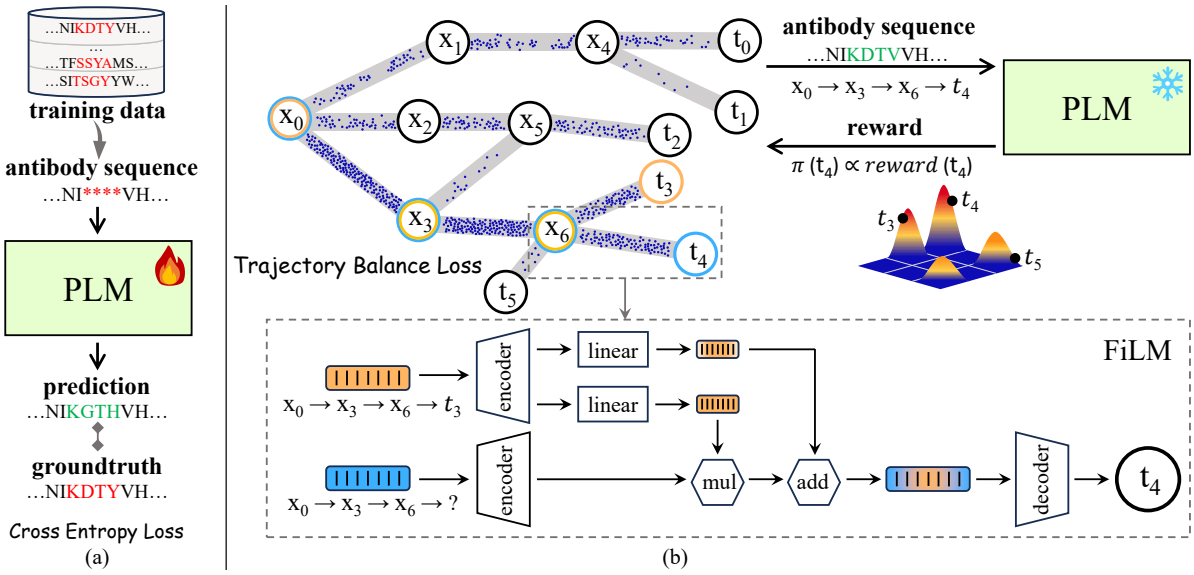

Figure 2: (a) Traditional approaches using pretrained PLM for antibody design. (b) Our PG-AbD framework that performs self-exploring towards pretrained PLM. Given pretrained PLM as the reward function, GFlowNet independently explores the rational design trajectories of antibody sequences. The FiLM module on the bottom exploits the template sequence to guide the internal exploration direction of GFlowNet. (Best viewed in color.)

GFlowNets, and trains the policy network with a fixed reward function. Additionally, intrinsic rewards could be introduced to help exploration in sparse reward tasks, where the policy network is jointly trained with the reward function (Ren et al. 2022; Zhang et al. 2022). And the extensions of GFlowNets for variational EM (Hu et al. 2023), multi-objective optimization (Jain et al. 2023; Zhu et al. 2024), continuous spaces (Lahlou et al. 2023) and stochastic spaces (Pan et al. 2023b) have also been studied. GFlowNet is very appealing to many applications, owing to its design philosophy to generate a greater diversity of solutions. (Pan et al. 2023a) shows the efficacy of GFlowNets in molecule generation. (Jain et al. 2022; Ghari et al. 2023) exploit GFlowNets within an active learning loop to design biological sequences. GFlowNets have also been applied to robust combinatorial optimization (Zhang et al. 2023). But until now, there has been no work effectively applying GFlowNets for antibody CDR design.

2.3 Protein Language Models

Protein Language Models (PLMs) pretrained on large-scale protein sequence corpus acquire powerful representations based on evolutionary information. PLMs have shown great promise in numerous protein downstream tasks, including function classification (Elnaggar et al. 2022; Yin et al. 2024), mutation effects prediction (Meier et al. 2021), structure prediction (Lin et al. 2023), and designing diverse and novel protein sequences (Nijkamp et al. 2023; Olsen, Moal, and Deane 2022; Melnyk et al. 2023). But directly exploiting pretrained PLMs for antibody CDR design often fails to capture the extensive variety of plausible sequences, limiting the generation diversity and novelty of feasible solutions.

3 Methodology

3.1 Preliminaries

We denote an antibody sequence as $s =< s_1, s_2, ..., s_n >$, where s_i represents the ith residue in an antibody sequence s of length n. Compared to Monte-Carlo Markov chains (MCMC), GFlowNet amortizes the expensive computational cost in a single but trained generative pass. GFlowNet models the generation process of antibody sequences by discrete actions, incrementally modifying a partially constructed object. The compositional objects are sampled sequentially, with each step involving the addition of a building block $a \in \mathcal{A}$ (action space) to the current partially constructed object $x \in \mathcal{X}$ (state space). A complete trajectory $\tau = (x_0 \rightarrow x_1 \rightarrow ... \rightarrow x_{n-1} \rightarrow x_n)$ is accomplished by a sequence of state transitions dictated by a series of actions. x_n is the terminal state, corresponding to a sample from the target space $\mathcal{S}$ of antibody sequences. The generative process of GFlowNet can be defined as a directed acyclic graph (DAG) $\mathcal{G} = (\mathcal{X}, \mathcal{A})$, with nodes and edges formed by partially constructed objects and action-decided states transitions, respectively. Note that, during the generative process of GFlowNet, multiple trajectories may lead to the same terminal state which provides GFlowNet detailed global probability modeling over the objective space. The forward policy $P_F(x'|x)$ is the collection of distributions over the children of states while the backward policy $P_B(x|x')$ is the collection of distributions over the parents of states. Given a non-negative reward function to score terminal state $R(x_n) : S \longmapsto \mathbb{R}^+$, the policy of GFlowNet is trained to sample objects with the probability proportional to reward ($\pi(s) \propto R(s)$). GFlowNet learns stochastic policies to sequentially generate compositional objects in proportion

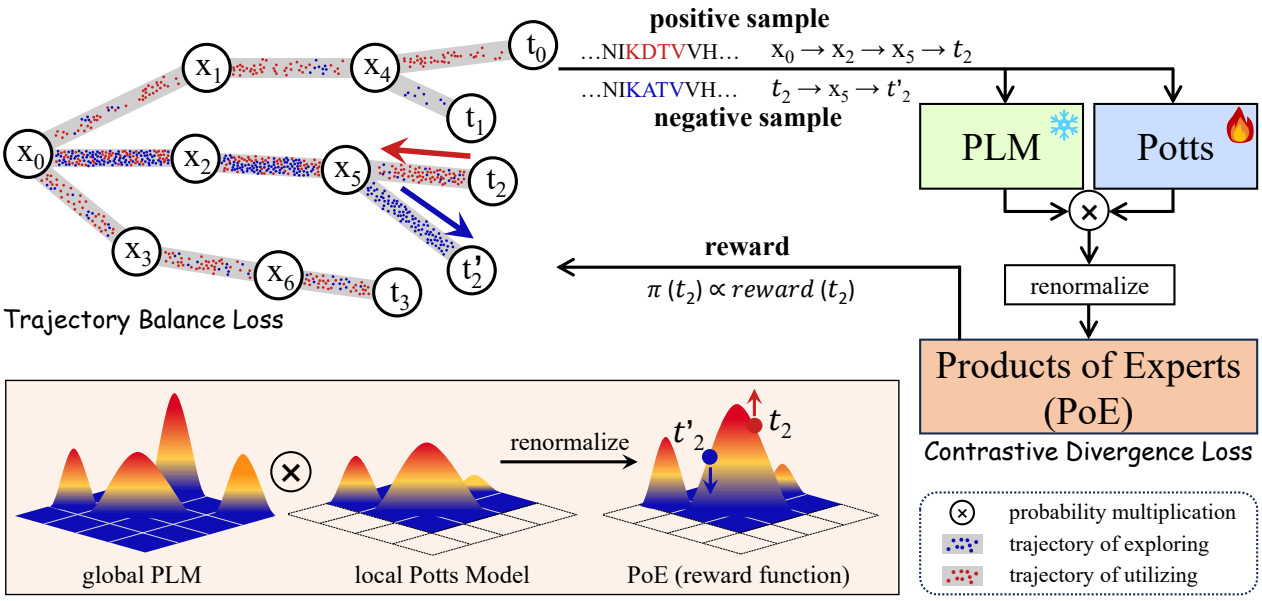

Figure 3: Our PG-AbD framework that jointly trains GFlowNet and PoE. GFlowNet provides positive-negative sample pairs for the contrastive learning of PoE, while PoE serves as the reward function to guide the exploration direction of GFlowNet. The constructed process of PoE is displayed on the bottom. Constraints in the global protein domain and local antibody specific domain are modeled by the pretrained PLM and a Potts Model, respectively. (Best viewed in color.)

to the given reward function and is well-suited to generate a diverse set of high-reward objects.

3.2 Overview

In this paper, we propose PG-AbD, which utilizes the synergy of GFlowNet and PLM for antibody design. Our framework is training-data-free. On the one hand, GFlowNet can operate in a self-exploring mode under the guidance of pretrained PLM, without training data. On the other hand, GFlowNet can execute both exploration and exploitation based on the training data trajectories, enabling the joint training with the PoE, constructed from the pretrained PLM and a Potts Model. Such unique property makes PG-AbD suitable not only for novel antibody design with limited training data, but also for antibody optimization when sufficient training sequences are available.

3.3 GFlowNet self-exploring towards PLM

GFlowNets generate antibody sequences $s \in \mathcal{S}$ with probability which is in proportion to a positive probability reward function $\pi(s) \propto R(s)$. The reward function of GFlowNet can be parameterized and its parameters need to be learned by maximum likelihood, which requires GFlowNet to sample from objective distribution. Moreover, the reward function can also take the form of a pretrained model acting as an external oracle to score samples derived by GFlowNet. In this paper, we leverage the pretrained auto-regressive protein language model ProGen2 (Nijkamp et al. 2023) as the reward function of GFlowNet, reflecting naturalness of the designed antibody sequences. As shown in Figure 2, the pretrained PLM with all-parameter fixed evaluates the naturalness of antibody sequences sampled by GFlowNet and guides GFlowNet to freely explore the rational generative trajectory of antibody sequences. Unlike typical neural networks, GFlowNet's training does not rely on external (real) data, it can take advantage of internally generated (pseudo) data, which means GFlowNet can automatically explore within the objective distribution without training data. Additionally, the FiLM module exploits the template sequence to guide the internal exploration direction of GFlowNet. Based on the feedback of constructed reward function, GFlowNet keeps on seeking and sampling rational antibody sequences. Ultimately, GFlowNet generates a diverse set of samples with high evaluated reward scores.

3.4 Joint training of PoE and GFlowNet

As described in Section 3.3, GFlowNet can independently explore the rational design trajectory of antibody sequences given a fixed PLM as reward function. However, the PLM pretrained on general protein data is not powerful enough to model antibody specific constraints. For comprehensively modeling constraints in both general protein domain and specific antibody domain, constructing a new reward function is worthwhile. Here, we further design a holistic reward function by exploiting a Products of Experts (PoE) constructed from the pretrained PLM and a Potts Model.

Potts Model The Potts Model, also known as Direct Coupling Analysis (DCA), has found numerous applications in protein biology, as a statistical model for the evolutionary process of protein sequences (Russ et al. 2020; Hopf et al. 2017; Trinquier et al. 2021). The Potts Model models pairwise constraints describing co-dependencies in combinations of amino acids for each pair of sites, and site-

specific constraints reflecting bias toward or away from specific amino acids at each position (Hopf et al. 2017).

$$P(s) = \frac{exp(\sum_i \mathbf{h}_i(s_i) + \sum_{i<j} \mathbf{J}_{ij}(s_i, s_j))}{Z}, \quad (1)$$

where $\mathbf{h}_i$ and $\mathbf{J}_{ij}$ represent site-wise bias terms and coupling terms, respectively. Z normalizes the distribution. Here, we exploit the Potts Model to explicitly model the antibody sequences probability distribution considering both the residue conservation and co-evolution.

Products of Experts The Products of Experts (PoE) combines multiple probabilistic models of the same data by multiplying their probability distributions together and then renormalizing (Hinton 1999, 2002). PoE acts as an effective and efficient probability function, where each individual expert model focuses on modeling distributions in a specific low-dimension space. Given global constraints learned by pretrained PLM in the general protein domain, we additionally introduce the Potts Model to model the local constraints in the specific antibody domain. As illustrated in Figure 3, we construct the PoE by multiplying the probabilities of global-distribution-modeling PLM and the local-distribution-modeling Potts Model together and then renormalizing. Serving as the reward function, PoE holistically models the probability distribution over antibody databases. Note that only samples simultaneously satisfying constraints in two distributional spaces (*i.e.*, global PLM and local Potts Model) will be assigned high probability.

Interleaved updating of PoE and GFlowNet Given the constructed PoE for comprehensive probabilistic modeling, now we can model the generative process under the guidance of PoE. Instead of the vanilla thoughts to sequentially train PoE and GFlowNet, we innovatively propose to jointly update constructed PoE and the policy network of GFlowNet motivated by (Zhang et al. 2022). With PoE acting as the intermediate reward function, GFlowNet is trained to convert the flow consistency equations, ensuring $\pi(s) \propto R(s)$. Updated GFlowNet also serves as a MCMC transition kernel to sample negative samples for PoE updating. Specifically, GFlowNet utilizes training data trajectories and performs back-and-forth transition proposal (Zhang et al. 2022; Kim et al. 2024). As illustrated in Figure 3, given the trajectories of antibody sequences in training database, we sample a K-step back trajectory with the backward policy $\mathbf{x}_i \sim P_B(\mathbf{x}_i|\mathbf{x}_{i+1})$:

$$\tau = (\mathbf{s} = \mathbf{x}_D \dashrightarrow \mathbf{x}_{D-1} \dashrightarrow \ldots \dashrightarrow \mathbf{x}_{D-K}). \quad (2)$$

Then we sample a forward trajectory with the forward policy to achieve a new terminal state, referring to as the negative sample of the antibody sequence $\mathbf{x}'_{i+1} \sim P_F(\mathbf{x}'_{i+1}|\mathbf{x}'_i)$:

$$\tau' = (\mathbf{x}'_{D-K} \rightarrow \ldots \rightarrow \mathbf{x}'_{D-1} \rightarrow \mathbf{x}'_D = \mathbf{s}'), \quad (3)$$

Similar to traditional MCMC methods, the Metropolis-Hastings (MH) rejection rule is applied to decide whether to accept the negative sample generated by GFlowNet.

$$A_{\tau,\tau'}(\mathbf{s} \rightarrow \mathbf{s}') \triangleq \min\left[1, \frac{e^{-\mathcal{E}_\phi(\mathbf{s}')}}{e^{-\mathcal{E}_\phi(\mathbf{s})}} \frac{P_B(\tau'|\mathbf{s}')P_F(\tau)}{P_B(\tau|\mathbf{s})P_F(\tau')}\right] \quad (4)$$

where $e^{-\mathcal{E}_\phi}$ refers to the constructed PoE for energy probabilistic modeling. We train PoE using the contrastive divergence, exploiting the positive-negative sample pairs generated by GFlowNet.

$$\mathbb{E}_{\mathbf{s}\sim p(\mathbf{s})}[\nabla_\phi \mathcal{E}_\phi(\mathbf{s}) - \mathbb{E}_{\mathbf{s}'\sim q_K(\mathbf{s}'|\mathbf{s})} \nabla_\phi \mathcal{E}_\phi(\mathbf{s}')] \quad (5)$$

Trained with this approximate MLE gradient, PoE becomes the comprehensive probability function, modeling both global constraints in the general protein domain and local constraints in the specific antibody domain. While PoE is gradually trained towards the likelihood of generating antibodies database, GFlowNet explores the generated trajectories of optimized antibody sequences under the guidance of updated PoE.

4 Experiments

We conduct three important yet challenging tasks to comprehensively evaluate the antibody generative capacity of PG-AbD framework: (1) Novel antibody Design with desired antigen binding modes (Section 4.1). (2) Antibody optimization by redesigning CDRs (Section 4.2). (3) Generated antibodies and antigens structural binding (Section 4.3).

CDR-H3						
w/o additional training data	Method	RMSD↓	TM↑	AAR↑	PPL↓	DIV↑
✓	RabD	-	-	28.53	8.72	-
✗	LSTM	-	-	21.97	8.07	-
	AR-GNN	3.60	-	19.74	8.00	20.06
	refine-GNN	**1.97**	**80.9**	34.99	6.54	27.87
✗	ProtBert	2.83	80.1	54.51	7.26	13.25
	AbLang-H	2.65	79.7	53.58	8.30	17.09
	AbLang-L	2.74	79.2	50.13	8,70	28.50
✓	ProtBert	2.78	78.6	33.29	5.11	55.29
	AbLang-H	3.43	78.4	6.23	13.96	-
	AbLang-L	3.77	78.1	4.64	10.52	-
	ESM-2	2.72	79.5	30.37	7.03	54.25
✗	reprogBERT	3.34	79.9	34.9	5.08	55.26
✓	PG-AbD (w/o FiLM)	2.63	80.1	35.17	5.06	**76.59**
✓	PG-AbD (w/ FiLM)	2.16	80.4	**58.40**	**4.26**	62.73

Table 1: Performance on RabDab. For pretrained PLMs, we simultaneously report the fine-tuned and zero-shot results.

Implement details. Our model is trained upon PyTorch framework using 4 NVIDIA GeForce RTX 4090 GPUs. For the deployment of model architecture, we exploit pretrained ProGen2-base (Nijkamp et al. 2023) with 764 million parameters for reward function, and a 4-layer transformer for policy network. We use learning rates of 1×10^{-4} and 5×10^{-3} to train PoE (*i.e.*, reward function) and policy network of GFlowNet respectively, with both being updated via the Adam optimizer. The code is available at https://github.com/KDurant-123/PG-AbD.

Baselines. We adopt LSTM from (Saka et al. 2021; Akbar et al. 2022), the energy-based model CEM (Fu and Sun 2022), and the diffusion probabilistic model AntibodyDiff (Luo et al. 2022). GNN-based antibody co-design methods, including AR-GNN, refine-GNN (Jin et al. 2021), are employed as baselines. Additionally, on the antigen-specific antibody design task, we include the physics-based

CDR-H1				
Method	AAR↑	PPL↓	NOV↑	DIV↑
LSTM†	-	6.1	-	-
AR-GNN†	-	6.0	-	-
refine-GNN	61.2	3.5	23.1	47.3
CEM†	39.8	6.1	-	5.3
AntibodyDiff†	41.2	6.9	-	23.5
ProtBert	64.7	3.5	18.0	4.6
AbLang-H	47.7	-	-	42.8
ESM-2†	50.1	4.2	29.9	37.7
reprogBERT	56.0	3.3	-	29.1
PG-AbD	**66.3**	**2.9**	**31.8**	**47.5**

CDR-H2				
Method	AAR↑	PPL↓	NOV↑	DIV↑
LSTM†	-	6.5	-	-
AR-GNN†	-	6.1	-	-
refine-GNN	48.9	**3.4**	28.6	38.7
CEM†	31.2	5.2	-	17.4
AntibodyDiff†	19.2	5.9	-	26.0
ProtBert	59.5	3.6	14.7	5.5
AbLang-H	46.7	-	-	44.9
ESM-2†	34.1	4.7	35.9	48.3
reprogBERT	53.0	3.9	-	37.9
PG-AbD	**61.2**	3.5	**37.0**	**49.1**

CDR-H3				
Method	AAR↑	PPL↓	NOV↑	DIV↑
LSTM†	-	8.2	-	-
AR-GNN†	-	8.5	-	-
refine-GNN	28.2	7.2	41.8	25.7
CEM†	23.7	7.9	-	26.1
AntibodyDiff†	37.6	9.0	-	23.9
ProtBert	**41.5**	6.8	33.7	14.5
AbLang-H	22.0	-	-	**71.3**
ESM-2†	32.4	5.9	35.9	33.4
reprogBERT	32.6	5.4	-	67.4
PG-AbD	39.5	**5.1**	**46.3**	67.8

Table 2: Performance on SabDab benchmark. †: re-implemented models from the released codes.

baseline called RosettaAntibodyDesign (RAbD) (Adolf-Bryfogle et al. 2018). Moreover, we focus on comparing with four performant PLMs (*i.e.*, ProtBert (Elnaggar et al. 2022), ESM-2 (Lin et al. 2023), AbLang (Olsen, Moal, and Deane 2022), and reprogBERT (Melnyk et al. 2023)).

Metrics. We adopt four metrics to thoroughly evaluate the developability, diversity and novelty of generated antibody sequences. we exploit perplexity (PPL) and amino acid recovery (AAR) for developability. Novelty (NOV) demonstrates the percent of generated antibodies with large amount of novel amino acids (*i.e.*, inconsistent with original amino acids). Diversity (DIV) computes the average recovery of all pairwise comparisons of the generated antibodies.

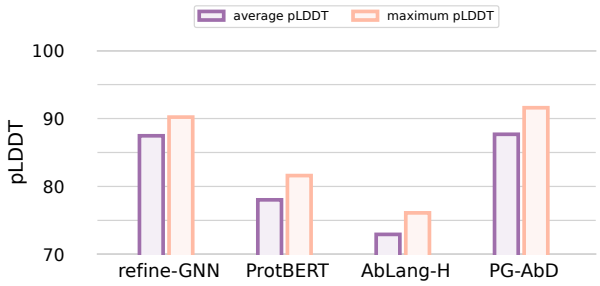

Figure 4: Results of Alphafold-Multimer structural confidence of generated antigen-antibody complexes.

4.1 Antigen-specific antibody design

Setup. This experiment is a real-world design task to generate antibodies binding to specific antigens based on limited antigen-antibody complex data. We select RabDab (Adolf-Bryfogle et al. 2018) as the benchmark (60 antibody sequences for test without training data). Most of existing approaches need additional training data (usually use antibody sequences from SabDab database (Dunbar et al. 2014)). In contrast, we exploit GFlowNet to perform self-exploring mode (introduced in Section 3.3) without requiring any additional training data.

Results. As shown in Table 1, our proposed framework provides convincing results of utilizing the pretrained PLM to generate diverse candidates for antibody discovery, simultaneously balancing the developability (high AAR and low PPL) and the diversity of the generated antibodies. In detail, PG-AbD obtains 62.73 DIV, which outperforms state-of-the-art methods by a large margin of 13.5%, while maintaining optimal 58.40 AAR and 4.26 PPL.

4.2 Antibody optimization

Setup. This experiment targets on the optimization of antibody CDRs from the SabDab (Dunbar et al. 2014). We jointly train PoE and GFlowNet (introduced in Section 3.4), where PoE models the antibody distribution probability over SabDab, and GFlowNet generates diverse CDR candidates for antibody optimization. The experimental results on SabDab are presented in Table 2.

Results. ProtBert generates candidates with high developability (high AAR and low PPL), but lacks diversity and novelty. AbLang and ESM-2 generate antibodies with high diversity and novelty, but seriously suffers from low developability. In contrast, our PG-AbD framework balances the developability, diversity and novelty of the generated antibodies and improves reprogBERT results by a large margin, leading to 18.3% AAR and 31.1% DIV gains on average.

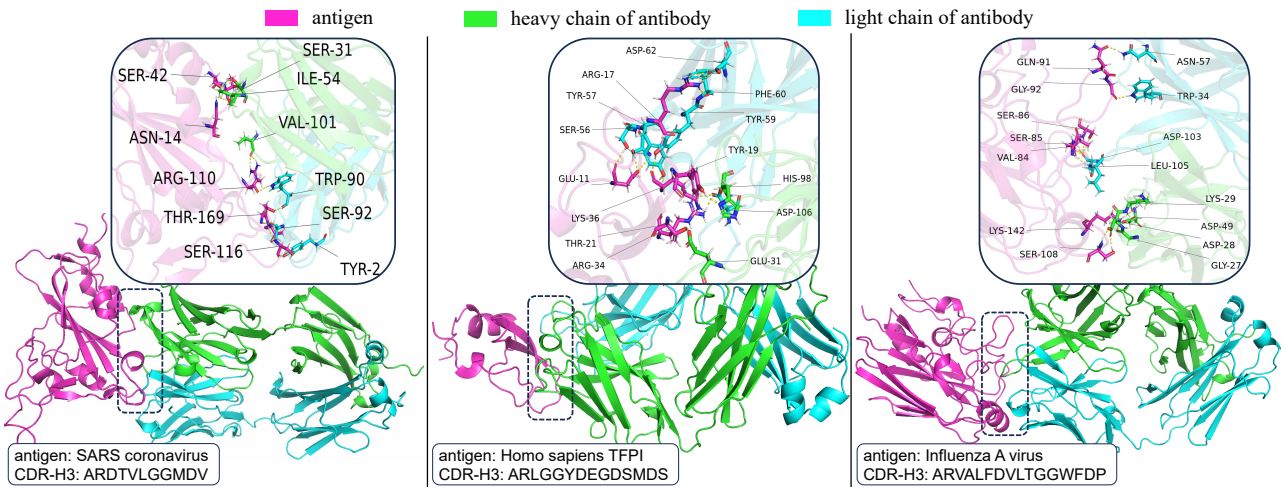

Figure 5: The structures of designed complexes, formed by the binding of PG-AbD generated antibodies and corresponding antigens. The detailed structures of antigen-antibody binding interfaces are visualized in the right top corner.

4.3 Structural binding of generated complex

Setup. In order to further evaluate the structural rationality and functionality of generated antibodies, complex structural binding experiments are incorporated to evaluate the binding effectiveness. Specifically, we collect a diverse complex dataset selected from (Adolf-Bryfogle et al. 2018). For each antigen-antibody complex, 100 CDR-H3 samples are generated. We select the top 10 antibody heavy chains with the minimum perplexity. Then Alphafold-Multimer (Evans et al. 2021) is used to predict the 3D structure of complexes comprising antigen and generated antibodies (including light chains and selected heavy chains), and record the average and maximum pLDDT values.

Results. As shown in Figure 4, PG-AbD generates complex structures with extremely higher confidence compared to traditional PLM methods, achieving 85.68 average pLDDT and 91.59 maximum pLDDT. Additionally, our PG-AbD framework, which is an antibody sequences based framework, surprisingly achieves the same level of structural rationality compared to the co-design method refineGNN, which simultaneously considering antibody sequence and structure constraints during the training process. Furthermore, some successful designed complex cases against are illustrated in Figure 5.

4.4 Ablation study

Ablation Study on PLMs Selection To select the most suitable pretrained PLM serving as the reward function of GFlowNet, we incorporate five widely-used PLMs (*i.e.*, ESM-2, ProtBert, AbLang-H, reprogBERT, and ProGen2) for comparison. The experimental results are presented in Table 3, exploiting the experimental setting of Section 4.1.

Ablation Study on PoE Components We execute ablation study to evaluate the necessity of multiple components of the constructed PoE, exploiting the experimental setting of Section 4.2. The pretrained PLM models the global constraints in general protein domain, and the Potts Model models the local constraints in specific antibody domain. The experimental results are shown in Table 4. Overall, the performance will decay if any one of the components is absent.

PLMs	AAR↑	PPL↓	DIV↑
ESM-2 (Lin et al. 2023)	54.21	6.12	31.26
ProtBert (Elnaggar et al. 2022)	44.23	5.90	42.31
AbLang-H (Olsen, Moal, and Deane 2022)	21.93	6.32	52.11
reprogBERT (Melnyk et al. 2023)	19.74	8.91	54.42
ProGen2 (Nijkamp et al. 2023)	**58.40**	**4.26**	**62.73**

Table 3: Ablation results of various PLMs.

Method	AAR↑	PPL↓	NOV↑	DIV↑
PG-AbD	**55.67**	**3.83**	**38.37**	**54.80**
w/o Potts Model	46.38	4.48	32.51	39.26
w/o PLM	33.29	5.79	29.13	45.70

Table 4: Ablation results of multiple PoE components.

5 Conclusion

In this paper, we propose PG-AbD, demonstrating that the synergy of GFlowNet and protein language model makes a diverse antibody designer. We first exploit GFlowNet self-exploring under the guidance of pretrained PLM to accomplish novel antibody design. Then we construct a Products of Experts (PoE) composed by the global-distribution-modeling PLM and the local-distribution-modeling Potts Model. PG-AbD jointly trains GFlowNet and the PoE, successfully generating optimized antibodies. Extensive experiments show the effectiveness of PG-AbD on multiple tasks, significantly improving the diversity of generated antibody candidates by a large margin, while maintaining optimal developability and novelty. We believe that PG-AbD and follow-up methods may potentially accelerate antibody discovery projects in real-world scenarios.

Acknowledgments

This research was partially supported by National Natural Science Foundation of China under grants No. 12326612, Zhejiang Key R&D Program of China under grant No. 2024SSYS0026 and No. 2023C03053, Alibaba Research Intern Program.

References

Adolf-Bryfogle, J.; Kalyuzhniy, O.; Kubitz, M.; Weitzner, B. D.; Hu, X.; Adachi, Y.; Schief, W. R.; and Dunbrack, R. L. 2018. RosettaAntibodyDesign (RAbD): A general framework for computational antibody design. *Public Library of Science Computational Biology*, 14(4): e1006112.

Akbar, R.; Robert, P. A.; Weber, C. R.; Widrich, M.; Frank, R.; Pavlović, M.; Scheffer, L.; Chernigovskaya, M.; Snapkov, I.; Slabodkin, A.; Mehta, B. B.; Miho, E.; Lund-Johansen, F.; Andersen, J. T.; Hochreiter, S.; Hobæk Haff, I.; Klambauer, G.; Sandve, G. K.; and Greiff, V. 2022. In silico proof of principle of machine learning-based antibody design at unconstrained scale. *Monoclonal Antibodies in Immunodiagnosis and Immunotherapy*, 14(1): 2031482.

Bengio, E.; Jain, M.; Korablyov, M.; Precup, D.; and Bengio, Y. 2021. Flow Network based Generative Models for Non-Iterative Diverse Candidate Generation. In *Advances in Neural Information Processing Systems*.

Dunbar, J.; Krawczyk, K.; Leem, J.; Baker, T.; Fuchs, A.; Georges, G.; Shi, J.; and Deane, C. M. 2014. SAbDab: the structural antibody database. *Nucleic Acids Research*, 42(D1): D1140–D1146.

Elnaggar, A.; Heinzinger, M.; Dallago, C.; Rehawi, G.; Wang, Y.; Jones, L.; Gibbs, T.; Feher, T.; Angerer, C.; Steinegger, M.; Bhowmik, D.; and Rost, B. 2022. ProtTrans: Toward Understanding the Language of Life Through Self-Supervised Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10): 7112–7127.

Evans, R.; O'Neill, M.; Pritzel, A.; Antropova, N.; Senior, A.; Green, T.; Žídek, A.; Bates, R.; Blackwell, S.; Yim, J.; Ronneberger, O.; Bodenstein, S.; Zielinski, M.; Bridgland, A.; Potapenko, A.; Cowie, A.; Tunyasuvunakool, K.; Jain, R.; Clancy, E.; Kohli, P.; Jumper, J.; and Hassabis, D. 2021. Protein complex prediction with AlphaFold-Multimer. *bioRxiv*.

Fu, T.; and Sun, J. 2022. Antibody Complementarity Determining Regions (CDRs) design using Constrained Energy Model. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.

Gao, K.; Wu, L.; Zhu, J.; Peng, T.; Xia, Y.; He, L.; Xie, S.; Qin, T.; Liu, H.; He, K.; and Liu, T.-Y. 2023. Pre-training Antibody Language Models for Antigen-Specific Computational Antibody Design. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.

Ghari, P. M.; Tseng, A.; Eraslan, G.; Lopez, R.; Biancalani, T.; Scalia, G.; and Hajiramezanali, E. 2023. Generative flow networks assisted biological sequence editing. In *NeurIPS 2023 Generative AI and Biology (GenBio) Workshop*.

Hinton, G. E. 1999. Products of experts. *Proc Icann*, 1: 1–6.

Hinton, G. E. 2002. Training Products of Experts by Minimizing Contrastive Divergence. *Neural Computation*, 14(8): 1771–1800.

Hopf, T. A.; Ingraham, J. B.; Poelwijk, F. J.; Schärfe, C. P. I.; Springer, M.; Sander, C.; and Marks, D. S. 2017. Mutation effects predicted from sequence co-variation. *Nature Biotechnology*, 35(2): 128–135.

Hu, J. E.; Malkin, N.; Jain, M.; Everett, K. E.; Graikos, A.; and Bengio, Y. 2023. GFlowNet-EM for learning compositional latent variable models. In *International Conference on Machine Learning*.

Jain, M.; Bengio, E.; Garcia, A.-H.; Rector-Brooks, J.; Dossou, B. F. P.; Ekbote, C.; Fu, J.; Zhang, T.; Kilgour, M.; Zhang, D.; Simine, L.; Das, P.; and Bengio, Y. 2022. Biological Sequence Design with GFlowNets. In *International Conference on Machine Learning*.

Jain, M.; Raparthy, S. C.; Hernandez-Garcia, A.; Rector-Brooks, J.; Bengio, Y.; Miret, S.; and Bengio, E. 2023. Multi-Objective GFlowNets. In *International Conference on Machine Learning*.

Jin, W.; Barzilay, R.; and Jaakkola, T. 2022. Antibody-Antigen Docking and Design via Hierarchical Equivariant Refinement. In *International Conference on Machine Learning*.

Jin, W.; Wohlwend, J.; Barzilay, R.; and Jaakkola, T. 2021. Iterative Refinement Graph Neural Network for Antibody Sequence-Structure Co-design. *arXiv preprint arXiv:2110.04624*.

Kim, M.; Yun, T.; Bengio, E.; Zhang, D.; Bengio, Y.; Ahn, S.; and Park, J. 2024. Local Search GFlowNets. In *International Conference on Learning Representations*.

Kong, X.; Huang, W.; and Liu, Y. 2023. Conditional Antibody Design as 3D Equivariant Graph Translation. In *International Conference on Learning Representations*.

Lahlou, S.; Deleu, T.; Lemos, P.; Zhang, D.; Volokhova, A.; Hernández-García, A.; Ezzine, L. N.; Bengio, Y.; and Malkin, N. 2023. A theory of continuous generative flow networks. *arXiv preprint arXiv:2301.12594*.

Li, T.; Pantazes, R. J.; and Maranas, C. D. 2014. OptMAVEn – A New Framework for the de novo Design of Antibody Variable Region Models Targeting Specific Antigen Epitopes. *Public Library of Science One*, 9(8): e105954.

Lin, Z.; Akin, H.; Rao, R.; Hie, B.; Zhu, Z.; Lu, W.; Smetanin, N.; Verkuil, R.; Kabeli, O.; Shmueli, Y.; et al. 2023. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637): 1123–1130.

Luo, S.; Su, Y.; Peng, X.; Wang, S.; Peng, J.; and Ma, J. 2022. Antigen-Specific Antibody Design and Optimization with Diffusion-Based Generative Models for Protein Structures. In *Advances in Neural Information Processing Systems*.

Meier, J.; Rao, R.; Verkuil, R.; Liu, J.; Sercu, T.; and Rives, A. 2021. Language models enable zero-shot prediction of the effects of mutations on protein function. In *Advances in Neural Information Processing Systems*.

Melnyk, I.; Chenthamarakshan, V.; Chen, P.-Y.; Das, P.; Dhurandhar, A.; Padhi, I.; and Das, D. 2023. Reprogramming pretrained language models for antibody sequence infilling. In *International Conference on Machine Learning*.

Nijkamp, E.; Ruffolo, J. A.; Weinstein, E. N.; Naik, N.; and Madani, A. 2023. ProGen2: Exploring the boundaries of protein language models. *Cell Systems*, 14(11): 968–978.e3.

Olsen, T. H.; Moal, I. H.; and Deane, C. M. 2022. AbLang: an antibody language model for completing antibody sequences. *Bioinformatics Advances*, 2(1): vbac046.

Pan, L.; Malkin, N.; Zhang, D.; and Bengio, Y. 2023a. Better Training of GFlowNets with Local Credit and Incomplete Trajectories. *arXiv preprint arXiv:2302.01687*.

Pan, L.; Zhang, D.; Jain, M.; Huang, L.; and Bengio, Y. 2023b. Stochastic Generative Flow Networks. In *Uncertainty in Artificial Intelligence*.

Raybould, M. I. J.; Marks, C.; Krawczyk, K.; Taddese, B.; Nowak, J.; Lewis, A. P.; Bujotzek, A.; Shi, J.; and Deane, C. M. 2019. Five computational developability guidelines for therapeutic antibody profiling. *Proceedings of the National Academy of Sciences*, 116(10): 4025–4030.

Ren, Z.; Guo, R.; Zhou, Y.; and Peng, J. 2022. Learning Long-Term Reward Redistribution via Randomized Return Decomposition. In *International Conference on Learning Representations*.

Rives, A.; Meier, J.; Sercu, T.; Goyal, S.; Lin, Z.; Liu, J.; Guo, D.; Ott, M.; Zitnick, C. L.; Ma, J.; and Fergus, R. 2021. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. *Proceedings of the National Academy of Sciences*, 118(15): e2016239118.

Russ, W. P.; Figliuzzi, M.; Stocker, C.; Barrat-Charlaix, P.; Socolich, M.; Kast, P.; Hilvert, D.; Monasson, R.; Cocco, S.; Weigt, M.; and Ranganathan, R. 2020. An evolution-based model for designing chorismate mutase enzymes. *Science*, 369(6502): 440–445.

Saka, K.; Kakuzaki, T.; Metsugi, S.; Kashiwagi, D.; Yoshida, K.; Wada, M.; Tsunoda, H.; and Teramoto, R. 2021. Antibody design using LSTM based deep generative model from phage display library for affinity maturation. *Scientific Reports*, 11(1): 5852.

Shin, J.-E.; Riesselman, A. J.; Kollasch, A. W.; McMahon, C.; Simon, E.; Sander, C.; Manglik, A.; Kruse, A. C.; and Marks, D. S. 2021. Protein design and variant prediction using autoregressive generative models. *Nature Communications*, 12: 2403.

Trinquier, J.; Uguzzoni, G.; Pagnani, A.; Zamponi, F.; and Weigt, M. 2021. Efficient generative modeling of protein sequences using simple autoregressive models. *Nature Communications*, 12(1): 5800.

Vincent, D.; Ethan, P.; Nathan, S.; Florian, S.; de Vries, H.; Aaron, C.; and Yoshua, B. 2018. Feature-wise transformations.

Wang, D.; Ye, F.; and Zhou, H. 2023. On Pre-trained Language Models for Antibody. In *International Conference on Learning Representations*.

Xu, M.; Yuan, X.; Miret, S.; and Tang, J. 2023. ProtST: Multi-Modality Learning of Protein Sequences and Biomedical Texts. In *International Conference on Machine Learning*.

Yin, M.; Zhou, H.; Zhu, Y.; Lin, M.; Wu, Y.; Wu, J.; Xu, H.; Hsieh, C.-Y.; Hou, T.; Chen, J.; and Wu, J. 2024. Multi-Modal CLIP-Informed Protein Editing. *Health Data Science*, 4: 0211.

Zhang, D.; Dai, H.; Malkin, N.; Courville, A.; Bengio, Y.; and Pan, L. 2023. Let the Flows Tell: Solving Graph Combinatorial Optimization Problems with GFlowNets. In *Advances in Neural Information Processing Systems*.

Zhang, D.; Malkin, N.; Liu, Z.; Volokhova, A.; Courville, A.; and Bengio, Y. 2022. Generative Flow Networks for Discrete Probabilistic Modeling. In *International Conference on Machine Learning*.

Zhu, F.; Althaus, T.; Tan, C. W.; Costantini, A.; Chia, W. N.; Van Vinh Chau, N.; Van Tan, L.; Mattiuzzo, G.; Rose, N. J.; Voiglio, E.; and Wang, L.-F. 2022. WHO international standard for SARS-CoV-2 antibodies to determine markers of protection. *The Lancet. Microbe*, 3(2): e81–e82.

Zhu, Y.; Wu, J.; Hu, C.; Yan, J.; Hou, T.; Wu, J.; et al. 2024. Sample-efficient multi-objective molecular optimization with gflownets. In *Advances in Neural Information Processing Systems*.