
FRAGMENTGPT: A Unified GPT Model for Fragment Growing, Linking, and Merging in Molecular Design

Xuefeng Liu^{1,5*}, Songhao Jiang^{1*}, Qinan Huang², Tinson Xu³
Ian Foster^{1,5}, Mengdi Wang⁶, Hening Lin⁴, Rick Stevens^{1,5}

¹Department of Computer Science, University of Chicago

²Pritzker School of Molecular Engineering, University of Chicago

³Laboratory Schools, University of Chicago

⁴Department of Chemistry, University of Chicago

⁵Argonne National Laboratory

⁶AI Lab, Princeton University

Abstract

Fragment-Based Drug Discovery (FBDD) is a popular approach in early drug development, but designing effective linkers to combine disconnected molecular fragments into chemically and pharmacologically viable candidates remains challenging. Further complexity arises when fragments contain structural redundancies, like duplicate rings, which cannot be addressed by simply adding or removing atoms or bonds. To address these challenges in a unified framework, we introduce FRAGMENTGPT, which integrates two core components: (1) a novel chemically-aware, energy-based bond cleavage pre-training strategy that equips the GPT-based model with fragment growing, linking, and merging capabilities, and (2) a novel Reward Ranked Alignment with Expert Exploration (RAE) algorithm that combines expert imitation learning for diversity enhancement, data selection and augmentation for Pareto and composite score optimality, and Supervised Fine-Tuning (SFT) to align the learner policy with multi-objective goals. Conditioned on fragment pairs, FRAGMENTGPT generates linkers that connect diverse molecular subunits while simultaneously optimizing for multiple pharmaceutical goals. It also learns to resolve structural redundancies—such as duplicated fragments—through intelligent merging, enabling the synthesis of optimized molecules. FRAGMENTGPT facilitates controlled, goal-driven molecular assembly. Experiments and ablation studies on real-world cancer datasets demonstrate its ability to generate chemically valid, high-quality molecules tailored for downstream drug discovery tasks.

1 Introduction

Fragment-based drug design (FBDD) has become a mainstream strategy in early drug discovery, especially for challenging targets like protein–protein interactions, allosteric sites, and shallow or cryptic pockets [39]. Unlike high-throughput screening (HTS), which depends on large, complex compound libraries [18], FBDD employs small, low-molecular-weight fragments to probe binding surfaces with precision and efficiency [15, 18, 39, 29]. Though fragments typically bind weakly (micromolar–millimolar range), they offer high ligand efficiency and chemical tractability, making them ideal starting points for optimization [34]. In addition, fragment libraries, usually only hundreds to a few thousand compounds, cover chemical space more effectively than HTS due to the combinatorial potential of linking or growing fragments. Their small, hydrophilic nature often yields hits with

*Equal Contribution. Correspondence to: Xuefeng Liu <xuefeng@uchicago.edu>.

favorable solubility and reduced off-target risks. Moreover, sensitive biophysical methods such as NMR, X-ray crystallography, SPR, and thermal shift assays enable rapid screening and validation, making FBDD a powerful, efficient route for lead discovery.

The impact of FBDD is underscored by the approval of several clinically significant drugs. For example, Zelboraf (vemurafenib), a BRAF inhibitor for melanoma, was the first FDA-approved drug derived from FBDD [4]. It was followed by others such as Venetoclax, a Bcl-2 inhibitor disrupting protein–protein interactions [7], and Sotorasib, a KRAS-G12C inhibitor targeting a previously “undruggable” oncogene [17]. In total, eight FDA-approved drugs and over 50 clinical candidates have originated from FBDD pipelines, validating its versatility across diverse target classes, including kinases, protein–protein interactions, and even RNA-binding proteins [39].

Despite its success, FBDD faces notable challenges. Once a fragment hit is identified, a major challenge lies in transforming these minimal binders into high-affinity ligands while maintaining or enhancing their binding specificity and desirable physicochemical traits. This is generally accomplished through three primary optimization strategies: fragment growing, fragment merging, and fragment linking.

Fragment growing involves expanding a single fragment into nearby binding regions, guided by structural information or structure–activity relationships (SAR). The clinical success of fragment growing is exemplified by several notable drug discoveries. Smith and colleagues advanced MRTX1719, targeting the PRMT5-MTA complex, beginning with 4-(aminomethyl)quinazoline-1(2H) identified through surface plasmon resonance [32]. Similarly, Heinrich and colleagues developed focal adhesion kinase inhibitors starting from 1H-pyrrolo[2,3-b]pyridine scaffolds [12], while Saxty et al. utilized their “Pyramid” X-ray crystallography platform to evolve 5-methyl-4-phenyl-1H-pyrazole hits into potent protein kinase B inhibitors [38]. Likewise, Addie and colleagues at AstraZeneca started with a small hinge-binding pyrrolopyrimidine fragment and, using crystal structures to guide each tweak, built it into AZD5363 (capiwasertib)—an oral, ATP-competitive Akt inhibitor with broad isoform potency, better ROCK selectivity, low hERG liability, and strong anti-tumor activity in cells and animals [2]. Fragment growing offers distinct advantages including efficient exploration of chemical space through incremental expansion, structure-guided optimization based on detailed binding insights, and typically higher ligand efficiency compared to traditional high-throughput screening approaches. The method enables rational design decisions by leveraging crystallographic or NMR structural information to direct growth toward productive binding site regions.

However, fragment growing presents significant challenges that must be carefully addressed. Primary considerations include determining optimal growth directions to maximize binding affinity while avoiding steric conflicts, maintaining geometric compatibility between added fragments and binding sites, and preserving drug-like properties throughout the expansion process. Additionally, ensuring synthetic accessibility of designed compounds remains crucial for practical development timelines and costs. The success of fragment growing depends on balancing multiple competing objectives: enhancing binding affinity and selectivity while maintaining favorable pharmacokinetic properties, drug-likeness, and synthetic feasibility. This multi-parameter optimization requires integration of structural biology, medicinal chemistry expertise, and sophisticated computational design tools to navigate the complex structure–activity relationships that emerge as fragments evolve into lead compounds.

Fragment merging integrates two or more overlapping fragments with shared substructures or binding features to preserve favorable interactions and enhance potency. Traditionally, this requires crystal structures to pinpoint overlaps—a labor-intensive process demanding expert knowledge—and even then, maintaining bioactive conformations often necessitates extensive optimization. Successful examples include trypanothione reductase inhibitors, where a propylphenyl and p-fluorophenyl group were merged via a piperazine scaffold [6], and HSP90 inhibitors, where adjacent fragments sharing a phenol moiety were combined but initially showed weak activity due to desolvation penalties [30]. Computational tools like the Fragment Network extend this idea by searching over 120 million catalogue compounds to suggest merges beyond simple fingerprint similarity [36], offering access to new regions of chemical space. However, such methods remain constrained by database coverage and lack true generative flexibility, leaving fragment merging an appealing yet underdeveloped strategy in practice.

Fragment linking connects two distinct fragments, each binding in adjacent pockets, using a chemical linker. Although both merging and linking offer conceptual advantages, they are notoriously

challenging to implement: merging demands precise alignment of shared pharmacophores and often relies on high-resolution co-crystal structures, while linking must address geometric limitations and design appropriate linkers to avoid entropic losses and ensure both fragments contribute to binding. These strategies frequently necessitate iterative synthesis guided by expert intuition. The increasing size of fragment libraries and growing demand for rapid optimization further highlight the need for computational tools—including virtual screening and AI-driven design—to enhance hit identification and linker optimization.

Recent advances in machine learning have introduced generative models as powerful tools for automated linker design. Sequence-based models, such as SyntaLinker [41], formulate linker generation as a sequence-to-sequence task, treating linkers as SMILES strings and employing natural language processing techniques, particularly Transformer-based neural networks. These models offer fast inference and strong chemical priors, effectively capturing the syntactic and semantic regularities of chemical representations. However, their reliance on 1D chemical language often necessitates post-processing steps to ensure the 3D structural feasibility and compatibility of generated linkers with protein binding pockets. In contrast, 3D geometry-based models, including those based on diffusion processes and variational autoencoders (VAEs)—such as 3DLinker [13], LinkerNet [10], DiffLinker [14], and DiffPROTACs [19]—generate linkers directly in three-dimensional space. These models exploit E(3)-equivariant architectures that inherently preserve rotational and translational symmetries, enabling the generation of spatially compatible and structurally plausible linkers conditioned on the geometric context of fragment binding poses and protein pockets.

In particular, DiffLinker, a state-of-the-art E(3)-equivariant 3D-conditional diffusion model, has demonstrated impressive performance in generating chemically valid, diverse, and pocket-compatible linkers between arbitrary numbers of fragments [14]. By modeling atomic geometries directly and conditioning on both fragment coordinates and protein pockets, DiffLinker surpasses prior approaches based on SMILES or autoregressive graph construction in terms of synthetic accessibility, drug-likeness, and similarity of samples to reference compounds. However, despite these achievements, current methods—including DiffLinker—focus solely on linking distinct fragments and remain limited in handling fragment merging, where two overlapping or partially redundant fragments must be merged into a single coherent substructure. No existing generative framework can yet simultaneously address both the linking and merging tasks in fragment-based design. Developing such unified strategies is critical for enabling more flexible and realistic drug design workflows where overlapping binding poses or scaffold overlaps frequently occur.

Currently, no generative framework can jointly address both fragment linking and merging and growing. The rise of large language models (LLMs) introduces a powerful paradigm for generative chemistry, yet their potential in fragment-based design—particularly linking and merging—remains largely underexplored. Moreover, pretrained GPT models rely solely on maximum likelihood training, lacking explicit objective-driven optimization. What’s more, pretrained GPT models have shown encouraging results in molecular tasks, highlighting the opportunity to harness these expert models [22, 24] to boost the performance of our framework. Finally, biomolecular design is inherently a multi-objective optimization problem. Multi-property evaluation measures how well a model balances multiple, often conflicting, objectives at once. One approach is a composite score, where different property metrics (e.g., LogP, QED, synthetic accessibility) are combined into a single weighted objective, guiding optimization methods. Another is Pareto optimality, where solutions are considered optimal if no property can be improved without worsening another; the collection of these solutions forms the Pareto front, which helps visualize trade-offs among objectives. However, both approaches come with their own advantages and limitations. To address these challenges, we asked

How can language models be employed to achieve fragment growing, linking, and merging, while integrating chemically-aware and expert knowledge to improve diversity under Pareto and composite score optimality alignment?

In this work, we introduce the first GPT-based framework that unifies fragment growing, linking, and merging within a single language model. Our method incorporates a novel pretraining strategy infused with chemically-aware and energy-based knowledge, allowing the model to execute these tasks without relying on pre-existing fragment databases. Instead, it directly identifies linkers and mergeable substructures from the molecular context. To further guide generation, we propose a Reward Ranked Alignment with Expert Exploration (RAE) algorithm that combines expert imitation learning with supervised fine-tuning. Expert imitation learning, leveraging pretrained expert models, promotes

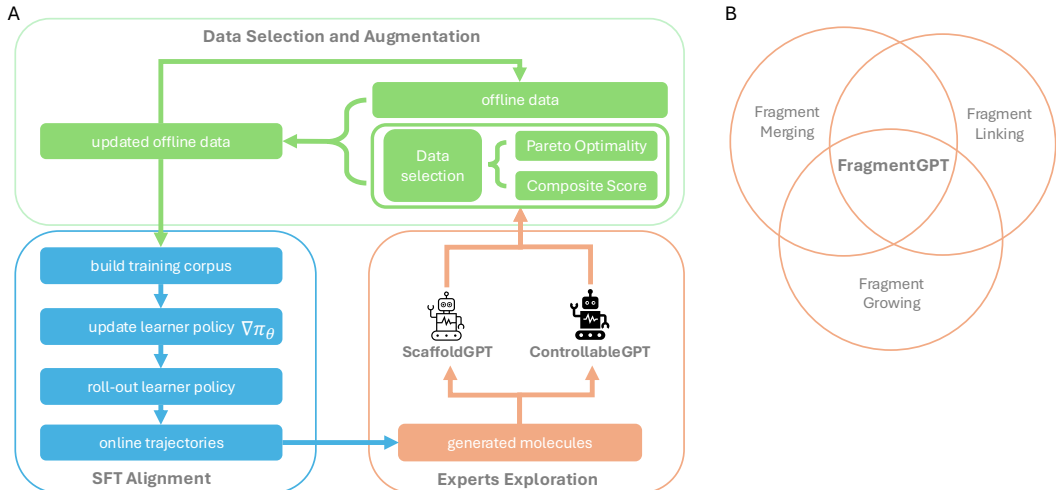


Figure 1: (A) The framework of Reward Ranked Alignment with Expert Exploration (RAE). (B) Functional overview of FRAGMENTGPT, which supports fragment merging, linking, and growing.

diversity, while supervised fine-tuning—together with data selection and augmentation—drives the learner policy toward Pareto-optimal and composite score solutions in multi-objective molecular property optimization. Taken together, the integration of GPT-driven fragment pretraining, expert-guided imitation, and targeted data augmentation establishes a scalable, adaptable, and goal-oriented framework that surpasses conventional database-driven approaches.

2 Preliminaries

LM/CLM. Let $F = [f_1, f_2, \dots, f_n]$ denote an input sequence of tokens (fragment-based prompt), where each f_i belongs to a vocabulary \mathcal{V} . The output sequence is defined as $M = [m_1, m_2, \dots, m_T]$ with $m_i \in \mathcal{M}$, where \mathcal{V} and \mathcal{M} may be distinct vocabularies. For convenience, we write $\mathbf{m}_{<t} = [m_1, \dots, m_{t-1}]$ and $\mathbf{m}_T := M$, and use T to denote sequence length. Each training sample begins with a start token [BOS], followed by a sequence \mathbf{m} with $y_i \in \mathcal{V}$, and terminates with an end token [EOS]. Molecules are represented as token sequences \mathbf{m} that assemble into SMILES strings, covering both partial and complete structures. We denote string concatenation by \circ and the Kleene closure of \mathcal{V} by \mathcal{V}^* . The training corpus C is thus: $C := \{[\text{BOS}] \circ \mathbf{v} \circ [\text{EOS}] \mid \mathbf{v} \in \mathcal{V}^*\}$. The generator policy π_θ , parameterized by a deep neural network with weights θ , defines a distribution over sequences as $\pi_\theta(\mathbf{m}|\mathbf{f}) = \prod_{t=1}^{|\mathbf{m}|} \pi_\theta(m_t|\mathbf{f}, \mathbf{m}_{<t})$, where $\pi_\theta(m_t|\mathbf{f}, \mathbf{m}_{<t}) = P(m_t|\mathbf{m}_{<t}, X)$ denotes the conditional probability of token m_t . Decoding aims to select the most likely sequence from the hypothesis space: $\mathbf{m}^* = \arg \max_{\mathbf{m} \in \mathcal{Y}_T} \log \pi_\theta(\mathbf{m}|\mathbf{f})$. Causal language modeling (CLM) reformulates the training objective as $\max_{\theta} \sum_{i=1}^n \log P(f_i|F_{<i}; \theta)$, where $P(f_i|F_{<i}; \theta)$ is the conditional probability of token f_i given its history. In this work, we employ a GPT policy π_θ trained on dataset \mathcal{B} to capture prior chemical knowledge and generate chemically valid molecules. Molecules are represented using SMILES strings [35]. The tokenizer is first pre-trained with the Byte Pair Encoding (BPE) algorithm [9, 31]. Building on this, we adopt an incremental pre-training strategy inspired by ScaffoldGPT [26], which significantly enhances the model’s ability to reconstruct valid molecules from fragments.

Fragment-based Drug Discovery. Given an initial prompt (fragment candidates) $\{F_{i=1}^n\} \sim \mathcal{X}$ and a drug discovery policy π_θ , the goal in fragment-based drug discovery is to find the optimal policy π_{θ^*} that maximizes the following objective:

$$\pi_{\theta^*} = \arg \max_{\pi_\theta} \mathbb{E}_{\{F\} \sim d_0} [R(M)|\theta, \{F_{i=1}^n\}], \quad (1)$$

where $M = \pi_\theta(\cdot | \{F_{i=1}^n\})$, $M_{1:T} = (m_1, \dots, m_t, \dots, m_T)$, $m_t \in \mathcal{V}$. Regards to $\{F_{i=1}^n\}$, For $\{F_{i=1}^n\}$, F_1 corresponds to fragment growing, F_1^2 corresponds to fragment linking and merging, and F_1^n with $n > 2$ corresponds to the linking and merging of multiple fragments.




Figure 2: Overview of training corpus construction for linking, merging and growing. For linking and merging, given a SMILES string, we divide it into three parts, which we denote as fragments *A*, *B*, and *C*. For the construction of the *linker training corpus*, we use fragments *A* and *C* as inputs, and represent each training instance as $\langle p1 \rangle \text{SMILES}_A \langle p2 \rangle \text{SMILES}_C \langle L \rangle \text{SMILES}_{ABC}$. Similarly, for the *merging training corpus*, we take the concatenated fragments *AB* and *BC* as inputs, resulting in the representation $\langle p1 \rangle \text{SMILES}_{AB} \langle p2 \rangle \text{SMILES}_{BC} \langle L \rangle \text{SMILES}_{ABC}$. For growing given fragment *A* and a molecule *M* which contains *A* as substructure, the model will take $\langle p1 \rangle \text{SMILES}_A \langle L \rangle \text{SMILES}_M$

3 The FRAGMENTGPT Algorithm

In this work, we bridge the gap among fragment growing, linking, and merging by introducing FRAGMENTGPT. FRAGMENTGPT directly trains a GPT model on SMILES representations, enabling it to generate complete molecules from fragment-based prompts. Furthermore, to address multi-objective molecular optimization, we propose RAE (Reward-Ranked Alignment with Expert Exploration)—a framework that combines supervised fine-tuning (SFT) retraining with expert-guided exploration, data selection, and augmentation.

Stage 1. Reward Ranked Alignment.

Pre-Training Corpus Construction. To enable the model to perform three key operations—linking, merging, and growing—we explicitly construct pretraining corpus that reflects these tasks. Our approach begins by converting a database of molecules from their SMILES strings into graph representations, where atoms serve as nodes and bonds as edges. The core of our methodology is a chemically-aware, energy-based bond cleavage strategy. To ensure the generation of stable fragments, we annotate each bond with its dissociation energy and explicitly prevent the cleavage of bonds within aromatic rings. The algorithm preferentially targets weaker bonds for disconnection, guided by two key parameters: a maximum of three fragments per molecule to ensure substructures are significant, and a bond energy cutoff of 90 kcal/mol (377 kJ/mol) to selectively cleave common single bonds while preserving stronger ones. By applying this graph-based fragmentation algorithm across a large-scale molecular database, we have successfully generated a comprehensive corpus. This dataset is consist of logically paired and chemically stable fragments.

Linking: After we constructed a large corpus of chemically meaningful fragments for linker and merger tasks, we generate a fragment-pair dataset. As illustrated in Figure 2, given a molecule *M*, we partition it into three components: *A*, *B*, *C*. This yields two classes of fragment pairs. First, the pair (*A*, *C*) where the task is to generate a linker bridging the two fragments. Each resulting fragment pair, together with the original molecule *M*, forms a training instance for linking or merging.

Merging: Second, the pairs (*AB*, *BC*) where the task is to merge fragments based on their shared substructure *B*. To further enrich the corpus and encourage the model to learn molecule merging, we construct additional examples by randomly selecting two molecules and merging them via their maximum common substructure (MCS). As shown in Figure 2, given two molecules MOL_A and MOL_B , we first

Algorithm 1 Reward Ranked Alignment with Expert Exploration (RAE)

Require: GPT-based generator policy π_θ^G ; offline dataset \mathcal{D} ; replay buffer \mathcal{B} ; number of experts E .

- 1: $\mathcal{D} \rightarrow \mathcal{B}$. Initialize π_θ^G with GPT2-like Transformer with random weight θ .
 - 2: **for** $n = 1, \dots, N$ **do**
 - ▷ /* Stage 1: SFT retraining */
 - 3: $\{M\}_{i=1}^N \sim \rho_0$, where $\rho_0 \in \Delta(\mathcal{D})$. $\mathcal{B} \leftarrow \emptyset$.
 - 4: Build the prompt \mathcal{X} and training corpus $C_{\text{growing}}, C_{\text{linking}}, C_{\text{merging}}$ (2) by using $\{M\}_{i=1}^N$.
 - 5: Pre-train BPE tokenizer and π_θ on $C_{\text{growing}}, C_{\text{linking}}$ and C_{merging} via CLM objective (3).
 - 6: Generate $M_{i=1}^N = \{(m_t, \dots, m_T)\}_{i=1}^N \sim \pi_\theta(\cdot | \{F_{i=1}^n\}), \{F_{i=1}^n\} \sim \mathcal{X}$.
 - 7: $\mathcal{B} \leftarrow \mathcal{B} \cup M_{i=1}^N$.
 - ▷ /* Stage 2: Expert Exploration */
 - 8: $\{M^{\text{EXPERT}}\}_{i=1}^N \sim \pi_j^{\text{EXPERT}}(M_i), i \in [N], j \in [E]$.
 - 9: $\mathcal{B} \leftarrow \mathcal{B} \cup \{M^{\text{EXPERT}}\}_{i=1}^N$.
 - ▷ /* Stage 3: Data Selection and Augmentation */
 - 10: $\mathcal{D} \leftarrow 80\% \cdot \text{PRIORITY}(\mathcal{B}, R^{\text{PARETO}}, R^{\text{COMPOSITE}}) \cup 20\% \cdot \text{RANDOM}(\mathcal{D})$.
-

identify their MCS (highlighted in red), then align and merge them to form a new molecule MOL_M . The corresponding training corpus is represented as $\langle p1 \rangle \text{SMILES}_A \langle p2 \rangle \text{SMILES}_B \langle L \rangle \text{SMILES}_M$ where the model learns to merge from two random molecules into a larger molecule.

Growing: Growing is used to further refine a fragment or a molecule obtained from linking or merging. The pretraining corpus is constructed as $\langle p1 \rangle \text{SMILES}_{A \text{ or } B} \langle L \rangle \text{SMILES}_M$, where SMILES may represent either a fragment or a complete molecule. The objective of growing is to transform the fragment or molecule into a more complete and refined structure.

The corpus C and prompts \mathcal{X} are constructed as follows:

$$\left\{ [BOS], \overbrace{\langle p1 \rangle, f_{11}, \dots, f_{1T}, \langle p2 \rangle, f_{21}, \dots, f_{2T}, \langle L \rangle}^{\text{PROMPT}}, \underbrace{m_1, \dots, m_T}_{\text{target molecule M}}, [EOS] \right\}. \quad (2)$$

Consequently, the model is able to generate the corresponding complete molecule when provided with a pair of fragments as the input prompt.

Pretrain a GPT generator. We denote the GPT generator policy by π_θ , which models the probability p of producing the t^{th} token in a target molecule Y . This probability is conditioned on all previously generated tokens $\mathbf{y}_{<t} = [y_1, \dots, y_{t-1}]$ in the sequence, as well as on the given pair of fragment compounds $\{F_{i=1}^n\}$, formally expressed as $\pi_\theta(y_t | \mathbf{y}_{<t}, \{F_{i=1}^n\}) = p(y_t | \mathbf{y}_{<t}, \{F_{i=1}^n\})$. The parameters θ of the generator policy π_θ are optimized by minimizing the negative log-likelihood (NLL) over the training corpus, computed on the complete SMILES sequences across the entire dataset. This process is described as follows:

$$NLL = -\log P(Y | \{F_{i=1}^n\}) \quad (3)$$

$$= -\sum_{t=1}^T \log P(y_t | y_{t-1}, \dots, y_1, \{F_{i=1}^n\}) = -\sum_{t=1}^T \log \pi_\theta(y_t | y_{1:t-1}, \{F_{i=1}^n\}), \quad (4)$$

Here, T denotes the total number of tokens in the target molecule Y . The negative log-likelihood (NLL) thus measures the likelihood of reconstructing a designated target molecule from the given fragment pair $\{F_{i=1}^n\}$. In this project, we adopt an incremental training strategy [26] that leverages large collections of unlabeled text to build a foundational model of language understanding. This foundation model can then be fine-tuned and adapted to a variety of downstream tasks with specific objectives.

Stage 2. Expert Exploration. We begin by sampling a batch of prompts $\mathcal{X}_t = \{\{F_{i=1}^n\}_1^t, \dots, \{F_{i=1}^n\}_b^t\} \subseteq \mathcal{X}$, and for each $\{F_{i=1}^n\}_i^t \in \mathcal{X}_t$, we generate molecules $M_1, \dots, M_K \sim \pi_\theta(\cdot | \{F_{i=1}^n\}_i^t)$. Next, we leverage the expert policy π^E to produce more diverse molecular candidates: $M_1^{\text{EXPERT}}, \dots, M_K^{\text{EXPERT}} \sim \pi^{\text{EXPERT}}(\cdot | M_i), i \in [K]$. Specifically, we

Target	Algorithm	Validity \uparrow	Avg Norm Reward \uparrow^*	Avg Top 10 % Norm Reward \uparrow	Docking \downarrow	Druglikeliness \uparrow	Synthesizability \downarrow	Solubility \uparrow	Similarity \uparrow
	Original	-	0.540	0.768	-8.505	0.720	2.981	2.364	-
Link	Link-INVENT [11]	0.517	0.515	0.609	-7.379	0.540	3.059	4.003	0.521
	DiffLinker [14]	0.891	0.428	0.655	-7.868	0.606	4.004	2.559	0.451
	FRAGMENTGPT (Pretrained)	0.979	0.563	0.738	-8.206	0.701	2.815	2.920	0.590
	FRAGMENTGPT (Finetuned)	0.991	0.616	0.734	-8.679	0.730	2.418	3.684	0.564
Merge	Link-INVENT [11]	0.409	0.443	0.510	-2.529	0.229	3.607	5.257	0.720
	DiffLinker [14]	0.858	0.345	0.550	-4.972	0.365	4.574	3.557	0.528
	FRAGMENTGPT (Pretrained)	0.947	0.630	0.788	-8.235	0.692	2.859	2.829	0.953
	FRAGMENTGPT (Finetuned)	0.993	0.651	0.791	-8.375	0.727	2.699	2.815	0.965

Table 1: **Main results.** A comparison of Link-INVENT, DiffLinker and different versions of FRAGMENTGPT on various objectives based on RTCB (PDBID: 4DWQ) datasets.

employ ScaffoldGPT [20, 26] and ControllableGPT [27, 21] to promote structural diversity. Finally, we combine the molecules generated by both the learner policy and the expert policy into the replay buffer \mathcal{B} , $\mathcal{B} \leftarrow \{M\} \cup \{M^{\text{EXPERT}}\}$.

Stage 3. Data Selection and Augmentation. In this step, we employ the Pareto optimality reward model $R^{\text{PARETO}}(M)$ from Abeer et al. [1] and the composite score reward model $R^{\text{COMPOSITE}}(M)$ from Liu et al. [26]. In this work, we regarded each pharmaceutical property as a critic C and considered an ensemble critics \mathbf{C} as follows:

$$[C^{\text{Druglikeliness}}, C^{\text{Solubility}}, C^{\text{Synthesizability}}, C^{\text{Docking}}, C^{\text{Tanimoto}}],$$

where each critic $C : M \rightarrow \mathbb{R}$ acts as a distinct evaluator for a specific pharmaceutical attribute. We built the composite score reward function as follow:

$$R^{\text{COMPOSITE}}(M) := \sum_{i=1}^{|\mathbf{C}|} \text{Norm}(C_i(M)),$$

where Norm is employed to standardize diverse attributes to a consistent scale. For Pareto optimality, we identify the Pareto front by iteratively collected fronts $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_n$ until the replay buffer \mathcal{B} is empty. These fronts partitions \mathcal{B} with $\mathcal{P}_1 \cap \mathcal{P}_2 = \emptyset$ for $i \neq j$ and $\mathcal{B} = \cup_i \mathcal{P}_i$. The ranking method assign each point $M \in \mathcal{P}_j$ the rank same as the reward as

$$R^{\text{PARETO}}(M) = \text{RANK}_{\mathcal{B}}(M) = \sum_i |\mathcal{P}_i|, \forall M \in \mathcal{P}_i. \quad (5)$$

Both models are used to evaluate each collected molecule in the replay buffer \mathcal{B} . We then reconstruct our offline dataset \mathcal{D}

$$\mathcal{D} \leftarrow 80\% \cdot \text{PRIORITY}(\mathcal{B}, R^{\text{PARETO}}, R^{\text{COMPOSITE}}) \cup 20\% \cdot \text{RANDOM}(\mathcal{D}) \quad (6)$$

by randomly sampling 20% from the previous \mathcal{D} and 80% from the buffer, with sampling prioritized according to their ranking [28]. We then proceed to Stage 1, where we retrain the learner policy on the dataset aligned with Pareto optimality and composite score.

4 EXPERIMENTS

4.1 Experimental configuration

The language model. We use GPT-2-like Transformers for causal language modeling. The training process is structured into three phases: pretraining, fine-tuning, as outlined in Algorithm FRAGMENT-GPT §3

Baselines. We benchmark our approach against baseline model, Link-INVENT [11], which employs an encoder-decoder architecture based on recurrent neural networks (RNNs). In this framework, two molecular fragments, referred to as *warheads*, are provided as input constraints to the RNN-based generator, which then produces a suitable linker to connect them. In addition, we include DiffLinker [14] in our comparison. DiffLinker employs an E(3)-equivariant three-dimensional conditional diffusion framework for molecular linker design, in which the SDF representations—chemical file formats that encode molecular structures along with their 3D coordinates and associated properties—of the two fragments are provided as input to guide the diffusion process.

Dataset. We utilize 1 million compounds from the ZINC15 database, docked against the human cancer protein RTCB (PDB ID: 4DWQ), as provided in the most recent Cancer dataset [23, 25].

Critics and evaluation metric. In this work, we assess the effectiveness of FRAGMENTGPT in generating molecules that exhibit favorable characteristics for pharmaceutical drug discovery. To perform this evaluation, we employ the RDKit chemoinformatics toolkit [16] and consider the following performance metrics. **Validity** determines whether the generated SMILES strings are syntactically correct. **Druglikeness** evaluates the probability that a molecule could serve as a viable drug candidate. **Solubility** estimates the likelihood that a molecule can dissolve in water, approximated by the water–octanol partition coefficient (LogP). **Synthesizability** measures the practical feasibility of chemical synthesis, where a score of 1 indicates high ease of synthesis and 10 indicates high difficulty [5]. **Docking Score** reflects the binding potential of a molecule to the target site; for computational efficiency, this score is predicted using a surrogate docking model (see Appendix A.4). **Similarity** is quantified by the Tanimoto coefficient, comparing the structural similarity between the generated SMILES and the original molecules. **Average Top 10% Normalized Reward** captures the mean normalized reward across the top 10% of molecules, ranked by their normalized reward. **Average Normalized Reward**, our primary evaluation metric, represents the mean normalized values of docking score, druglikeness, synthesizability, solubility, and similarity over all valid molecules. All evaluations are conducted on a sample of 1,280 molecules.

4.2 Main results

Table 1 demonstrates that FRAGMENTGPT substantially outperforms Link-INVENT across both linking and merging tasks. This advantage arises from both architectural and methodological differences: FRAGMENTGPT employs a transformer-based GPT backbone, which is more expressive than the RNN architecture of Link-INVENT, and it explicitly incorporates common substructures when merging fragments. In contrast, Link-INVENT neglects common substructures, often producing unnecessarily complex molecules when handling fragments with overlapping motifs, which leads to degraded performance. FRAGMENTGPT also significantly outperforms DiffLinker. For the separate task, FRAGMENTGPT achieves better performance due to its RAE fine-tuning, which enables it to generate higher-quality linkers and consequently more optimal resulting molecules.

For the merging task, a similar trend is observed: DiffLinker often produces unnecessarily complex molecules when dealing with fragments containing overlapping motifs, leading to degraded performance. Moreover, the finetuned version of FRAGMENTGPT consistently surpasses the pretrained model across nearly all metrics, highlighting the effectiveness of our proposed RAE algorithm in improving molecular generation.

Further experiments on *fragment growing* (Appendix A.1.1), *linking* (Appendix A.1.2), and *merging* (Appendix A.1.3) can be found in Appendix Appendix A.1.

5 CONCLUSION

We have introduced FRAGMENTGPT, a novel GPT-based framework for fragment-based drug discovery that integrates chemically-aware, energy-based bond cleavage pre-training. This design equips the model with fragment growing, linking, and merging capabilities. In addition, we proposed Reward Ranked Alignment with Expert Exploration (RAE), a SFT alignment strategy that leverages expert imitation for diversity, data selection and augmentation for Pareto and composite score optimization, and supervised fine-tuning for aligning policy behavior with multi-objective drug design goals. Through comprehensive evaluations on real-world cancer-related datasets, FRAGMENTGPT consistently outperformed competitive baselines across a wide range of metrics, underscoring its effectiveness in generating drug-like molecules with improved pharmaceutical properties. While the current implementation is limited to SMILES-based representations, we are working to extend FRAGMENTGPT to broader molecular formats and structural representations.

Acknowledgements

We thank T. Brettin and the Argonne ACLF team for their support with computing resources. This work is supported in part by the RadBio-AI project (DE-AC02-06CH11357), U.S. Department of Energy Office of Science, Office of Biological and Environment Research, the Improve project under contract (75N91019F00134, 75N91019D00024, 89233218CNA000001, DE-AC02-06-CH11357, DE-AC52-07NA27344, DE-AC05-00OR22725), the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration.

References

- [1] ANM Nafiz Abeer, Nathan M Urban, M Ryan Weil, Francis J Alexander, and Byung-Jun Yoon. Multi-objective latent space optimization of generative molecular design models. *Patterns*, 5 (10), 2024. 7
- [2] Matt Addie, Peter Ballard, David Buttar, Claire Crafter, Gordon Currie, Barry R. Davies, Judit Debreczeni, Hannah Dry, Philippa Dudley, Ryan Greenwood, Paul D. Johnson, Jason G. Kettle, Clare Lane, Gillian Lamont, Andrew Leach, Richard W. A. Luke, Jeff Morris, Donald Ogilvie, Ken Page, Martin Pass, Stuart Pearson, and Linette Ruston. Discovery of 4-amino-n-[(1s)-1-(4-chlorophenyl)-3-hydroxypropyl]-1-(7h-pyrrolo[2,3-d]pyrimidin-4-yl)piperidine-4-carboxamide (azd5363), an orally bioavailable, potent inhibitor of akt kinases. *Journal of Medicinal Chemistry*, 56(5):2059–2073, 2013. doi: 10.1021/jm301762v. URL <https://doi.org/10.1021/jm301762v>. PMID: 23394218. 2
- [3] Melissa L Boby, Daren Fearon, Matteo Ferla, Mihajlo Filep, Lizbé Koekemoer, Matthew C Robinson, COVID Moonshot Consortium†, John D Chodera, Alpha A Lee, Nir London, et al. Open science discovery of potent noncovalent sars-cov-2 main protease inhibitors. *Science*, 382 (6671):eabo7201, 2023. 13
- [4] Gideon Bollag, Jack Tsai, Jian Zhang, Christina Zhang, Pamela Ibrahim, Keith Nolop, and Paul Hirth. Vemurafenib: the first drug approved for braf-mutant cancer. *Nature Reviews Drug Discovery*, 11(11):873–886, Nov 2012. doi: 10.1038/nrd3847. URL <https://doi.org/10.1038/nrd3847>. Epub 2012 Oct 12. PMID: 23060265. 2
- [5] Peter Ertl and Ansgar Schuffenhauer. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *Journal of cheminformatics*, 1:1–11, 2009. 8
- [6] Cécile Exertier, Alessandra Salerno, Lorenzo Antonelli, Annarita Fiorillo, Riccardo Ocello, Francesca Seghetti, Jessica Caciolla, Elisa Uliassi, Matteo Masetti, Eleonora Fiorentino, Stefania Orsini, Trentina Di Muccio, Andrea Ilari, and Maria Laura Bolognesi. Fragment merging, growing, and linking identify new trypanothione reductase inhibitors for leishmaniasis. *Journal of Medicinal Chemistry*, 67(1):402–419, 2024. doi: 10.1021/acs.jmedchem.3c01439. URL <https://doi.org/10.1021/acs.jmedchem.3c01439>. PMID: 38164929. 2
- [7] W. J. Fairbrother, J. D. Levenson, D. Sampath, and A. J. Souers. Discovery and development of venetoclax, a selective antagonist of bcl-2. In J. Fischer, C. Klein, and W. E. Childers, editors, *Successful Drug Discovery*, volume 4, pages 225–245. Wiley-VCH, Weinheim, 2019. [Google Scholar]. 2
- [8] Leonardo G. Ferreira and Adriano D. Andricopulo. From protein structure to small-molecules: recent advances and applications to fragment-based drug discovery. *Current topics in medicinal chemistry*, 17(20):2260–2270, 2017. 13
- [9] Philip Gage. A new algorithm for data compression. *The C Users Journal*, 12(2):23–38, 1994. 4
- [10] Jiaqi Guan, Xingang Peng, Peiqi Jiang, Yunan Luo, Jian Peng, and Jianzhu Ma. Linkernet: fragment poses and linker co-design with 3d equivariant diffusion. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS ’23, Red Hook, NY, USA, 2023. Curran Associates Inc. 3

- [11] Jeff Guo, Franziska Knuth, Christian Margreitter, Jon Paul Janet, Kostas Papadopoulos, Ola Engkvist, and Atanas Patronov. Link-invent: generative linker design with reinforcement learning. *Digital Discovery*, 2(2):392–408, April 2023. ISSN 2635-098X. doi: 10.1039/D2DD00115B. URL <https://pubs.rsc.org/en/content/articlelanding/2023/dd/d2dd00115b>. 7
- [12] Timo Heinrich, Jeyaprakashnarayanan Seenisamy, Lourdusamy Emmanuvel, Santosh Kulkarni, Jorg Bomke, Felix Rohdich, Hartmut Greiner, Christina Esdar, Mireille Krier, Ulrich Gradler, et al. Fragment-based discovery of new highly substituted 1 h-pyrrolo [2, 3-b]- and 3 h-imidazolo [4, 5-b]-pyridines as focal adhesion kinase inhibitors. *Journal of medicinal chemistry*, 56(3): 1160–1170, 2013. 2
- [13] Yinan Huang, Xingang Peng, Jianzhu Ma, and Muhan Zhang. 3DLinker: An e(3) equivariant variational autoencoder for molecular linker design. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 9280–9294. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/huang22g.html>. 3
- [14] Ilia Igashov, Hannes Stärk, Clément Vignac, Arne Schneuing, Victor Garcia Satorras, Pascal Frossard, Max Welling, Michael Bronstein, and Bruno Correia. Equivariant 3d-conditional diffusion model for molecular linker design. *Nature Machine Intelligence*, 6(4):417–427, April 2024. ISSN 2522-5839. doi: 10.1038/s42256-024-00815-9. URL <https://www.nature.com/articles/s42256-024-00815-9>. 3, 7
- [15] György M. Keserű, Daniel A. Erlanson, György G. Ferenczy, Michael M. Hann, Christopher W. Murray, and Stephen D. Pickett. Design principles for fragment libraries: Maximizing the value of learnings from pharma fragment-based drug discovery (fbdd) programs for use in academia. *Journal of Medicinal Chemistry*, 59(18):8189–8206, 2016. doi: 10.1021/acs.jmedchem.6b00197. URL <https://doi.org/10.1021/acs.jmedchem.6b00197>. PMID: 27124799. 1
- [16] Greg Landrum et al. RDkit: Open-source cheminformatics software. <https://www.rdkit.org>. Accessed Oct 2023. 8
- [17] Brian A. Lanman, Jennifer R. Allen, John G. Allen, Albert K. Amegadzie, Kate S. Ashton, Shon K. Booker, Jian Jeffrey Chen, Ning Chen, Michael J. Frohn, Guy Goodman, David J. Kopecky, Longbin Liu, Patricia Lopez, Jonathan D. Low, Vu Ma, Ana E. Minatti, Thomas T. Nguyen, Nobuko Nishimura, Alexander J. Pickrell, Anthony B. Reed, Youngsook Shin, Aaron C. Siegmund, Nuria A. Tamayo, Christopher M. Tegley, Mary C. Walton, Hui-Ling Wang, Ryan P. Wurz, May Xue, Kevin C. Yang, Pragathi Achanta, Michael D. Bartberger, Jude Canon, L. Steven Hollis, John D. McCarter, Christopher Mohr, Karen Rex, Anne Y. Saiki, Tisha San Miguel, Laurie P. Volak, Kevin H. Wang, Douglas A. Whittington, Stephan G. Zech, J. Russell Lipford, and Victor J. Cee. Discovery of a covalent inhibitor of krasg12c (amg 510) for the treatment of solid tumors. *Journal of Medicinal Chemistry*, 63(1):52–65, 2020. doi: 10.1021/acs.jmedchem.9b01180. URL <https://doi.org/10.1021/acs.jmedchem.9b01180>. PMID: 31820981. 2
- [18] Andrew R Leach and Michael M Hann. Molecular complexity and fragment-based drug discovery: ten years on. *Current Opinion in Chemical Biology*, 15(4):489–496, 2011. ISSN 1367-5931. doi: <https://doi.org/10.1016/j.cbpa.2011.05.008>. URL <https://www.sciencedirect.com/science/article/pii/S1367593111000718>. Next Generation Therapeutics. 1
- [19] Fenglei Li, Qiaoyu Hu, Yongqi Zhou, Hao Yang, and Fang Bai. Diffprotacs is a deep learning-based generator for proteolysis targeting chimeras. *Briefings in Bioinformatics*, 25(5):bbae358, 08 2024. ISSN 1477-4054. doi: 10.1093/bib/bbae358. URL <https://doi.org/10.1093/bib/bbae358>. 3
- [20] Xuefeng Liu, Songhao Jiang, Ian Foster, Jinbo Xu, and Rick L Stevens. Scaffold-driven gpt model for drug optimization. In *ICML 2025 Generative AI and Biology (GenBio) Workshop*, . 7
- [21] Xuefeng Liu, Songhao Jiang, Bo Li, and Rick L Stevens. A ground-up designed controllable gpt for molecule optimization. In *ICML 2025 Generative AI and Biology (GenBio) Workshop*, . 7

- [22] Xuefeng Liu, Fangfang Xia, Rick L Stevens, and Yuxin Chen. Cost-effective online contextual model selection. *arXiv preprint arXiv:2207.06030*, 2022. 3
- [23] Xuefeng Liu, Songhao Jiang, Archit Vasan, Alexander Brace, Ozan Gokdemir, Thomas Bretin, Fangfang Xia, Ian Foster, and Rick Stevens. Drugimprover: Utilizing reinforcement learning for multi-objective alignment in drug optimization. In *NeurIPS 2023 Workshop on New Frontiers of AI for Drug Discovery and Development*, 2023. 8, 14
- [24] Xuefeng Liu, Takuma Yoneda, Rick L Stevens, Matthew R Walter, and Yuxin Chen. Blending imitation and reinforcement learning for robust policy improvement. *arXiv preprint arXiv:2310.01737*, 2023. 3
- [25] Xuefeng Liu, Songhao Jiang, Siyu Chen, Zhuoran Yang, Yuxin Chen, Ian Foster, and Rick Stevens. Drugimprovergpt: A large language model for drug optimization with fine-tuning via structured policy optimization. *arXiv preprint arXiv:2502.07237*, 2025. 8
- [26] Xuefeng Liu, Songhao Jiang, Ian Foster, Jinbo Xu, and Rick Stevens. Scaffoldgpt: A scaffold-based gpt model for drug optimization. *arXiv preprint arXiv:2502.06891*, 2025. 4, 6, 7
- [27] Xuefeng Liu, Songhao Jiang, Bo Li, and Rick Stevens. Controllablegpt: A ground-up designed controllable gpt for molecule optimization. *arXiv preprint arXiv:2502.10631*, 2025. 7
- [28] Xuefeng Liu, Hung TC Le, Siyu Chen, Rick Stevens, Zhuoran Yang, Matthew R Walter, and Yuxin Chen. Active advantage-aligned online reinforcement learning with offline data. *arXiv preprint arXiv:2502.07937*, 2025. 7
- [29] Christopher W. Murray and David C. Rees. The rise of fragment-based drug discovery. *Nature Chemistry*, 1(3):187–192, 2009. doi: 10.1038/nchem.217. URL <https://doi.org/10.1038/nchem.217>. 1
- [30] Jie Ren, Jie Li, Yu Wang, Wei Chen, An Shen, Honglin Liu, Dan Chen, Dongmei Cao, Ying Li, Nan Zhang, Yili Xu, Meiyu Geng, Jian He, Bing Xiong, and Jianping Shen. Identification of a new series of potent diphenol hsp90 inhibitors by fragment merging and structure-based optimization. *Bioorganic & Medicinal Chemistry Letters*, 24(11):2525–2529, June 2014. ISSN 0960-894X. doi: 10.1016/j.bmcl.2014.03.100. URL <https://doi.org/10.1016/j.bmcl.2014.03.100>. Epub 2014 Apr 8. 2
- [31] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*, 2015. 4
- [32] Christopher R Smith, Ruth Aranda, Thomas P Bobinski, David M Briere, Aaron C Burns, James G Christensen, Jeffery Clarine, Lars D Engstrom, Robin J Gunn, Anthony Ivetac, et al. Fragment-based discovery of mrx1719, a synthetic lethal inhibitor of the prmt5• mta complex for the treatment of mta-deleted cancers. *Journal of Medicinal Chemistry*, 65(3):1749–1766, 2022. 2
- [33] Archit Vasan, Rick Stevens, Arvind Ramanathan, and Vishwanath Venkatram. Benchmarking language-based docking models. 2023. 16
- [34] Matthew J. Wasko, Kelly A. Pellegrine, Jeffery D. Madura, and Christopher K. Surratt. A role for fragment-based drug design in developing novel lead compounds for central nervous system targets. *Frontiers in Neurology*, 6:197, 2015. doi: 10.3389/fneur.2015.00197. URL <https://doi.org/10.3389/fneur.2015.00197>. 1
- [35] David Weininger. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of chemical information and computer sciences*, 28(1):31–36, 1988. 4
- [36] Stephanie Wills, Ruben Sanchez-Garcia, Tim Dudgeon, Stephen D. Roughley, Andy Merritt, Roderick E. Hubbard, James Davidson, Frank von Delft, and Charlotte M. Deane. Fragment merging using a graph database samples different catalogue space than similarity search. *Journal of Chemical Information and Modeling*, 63(11):3423–3437, 2023. doi: 10.1021/acs.jcim.3c00276. URL <https://doi.org/10.1021/acs.jcim.3c00276>. PMID: 37229647. 2

- [37] Georg E Winter, Andreas Mayer, Dennis L Buckley, Michael A Erb, Justine E Roderick, Sarah Vittori, Jaime M Reyes, Julia di Iulio, Amanda Souza, Christopher J Ott, et al. Bet bromodomain proteins function as master transcription elongation factors independent of cdk9 recruitment. *Molecular cell*, 67(1):5–18, 2017. [13](#)
- [38] Paul G Wyatt, Andrew J Woodhead, Valerio Berdini, John A Boulstridge, Maria G Carr, David M Cross, Deborah J Davis, Lindsay A Devine, Theresa R Early, Ruth E Feltell, et al. Identification of n-(4-piperidiny)-4-(2, 6-dichlorobenzoylamino)-1 h-pyrazole-3-carboxamide (at7519), a novel cyclin dependent kinase inhibitor using fragment-based x-ray crystallography and structure based drug design. *Journal of medicinal chemistry*, 51(16):4986–4999, 2008. [2](#)
- [39] Weijun Xu and Congbao Kang. Fragment-based drug design: From then until now, and toward the future. *Journal of Medicinal Chemistry*, 68(5):5000–5004, 2025. doi: 10.1021/acs.jmedchem.5c00424. URL <https://doi.org/10.1021/acs.jmedchem.5c00424>. PMID: 39992814. [1](#), [2](#)
- [40] Liu Xuefeng, Tien Chih-Chan, Ding Peng, Jiang Songhao, and Stevens Rick. Entropy-reinforced planning with large language models for de novo drug discovery. *arXiv 2024*, 2024. [16](#)
- [41] Yuyao Yang, Shuangjia Zheng, Shimin Su, Chao Zhao, Jun Xu, and Hongming Chen. Syn-talinker: automatic fragment linking with deep conditional transformer neural networks. *Chem. Sci.*, 11:8312–8322, 2020. doi: 10.1039/D0SC03126G. URL <http://dx.doi.org/10.1039/D0SC03126G>. [3](#)

A Appendix

A.1 Additional experiments

A.1.1 Fragment Growing

Figure 3 illustrates examples of fragment growing, where the original molecule was obtained from an NMR-based screening [8]. This example highlights the ability of FRAGMENTGPT to effectively transform simple fragments into chemically plausible and drug-like molecules.




Figure 3: Given an initial fragment, FRAGMENTGPT can effectively expand and refine it into a more complete molecular structure. The **highlighted substructures** indicate the newly grown parts.

A.1.2 Fragment Linking

dBET6, as visualized in figure 4, is a PROTAC molecule with high structural selectivity that bridges the E3 ubiquitin ligase CRBN and the target protein BRD4 [37]. We use this example to investigate whether our model can identify reasonable linkage sites and propose a suitable linker, and results are represented in figure 5.




Figure 4: dBET6 Chemical Structure

The molecular generation process predominantly yielded structures where the two fragments are connected via a linker, with only a few exceptions showing direct linking. All generated linkers were determined to be chemically stable. Notably, several of these linkers incorporate an imidazolinone scaffold. This structural motif is highly significant in medicinal chemistry, widely utilized in drug molecules for its chemical stability and biocompatibility. Furthermore, a particularly noteworthy achievement is that our model successfully identified the acyl linking site on one of the fragments, despite not being provided with any explicit information about this site in the inference stage. This demonstrates the model’s robust capability for implicit learning and recognition of key chemical features.

A.1.3 Fragment Merging

To assess our model’s ability to merge different fragments, we utilized the COVID Moonshot dataset [3]—a fully open science, crowd-sourced, and structure-enabled drug discovery campaign against the severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2) main protease (Mpro). This dataset provides an ideal testbed for evaluating fragment-based drug design approaches due to its comprehensive collection of crystallographic fragment hits and subsequent rational design submissions. The submission TRY-UNI-714a760b represents a class of rationally designed molecules that were constructed based on the strategic combination of five distinct fragments identified through crystallographic screening. These fragments were carefully selected and merged by human medicinal chemists who recognized their spatial complementarity and potential for synergistic binding interactions within the Mpro active site. To evaluate our model’s fragment merging capabilities, we challenged it to




Figure 5: Chemical structure of a bifunctional PROTAC degrader. The molecule is designed to induce proximity between an E3 ubiquitin ligase and a Protein of Interest (POI). The E3-recruiting fragment is colored red, the POI-binding warhead is blue, and the interconnecting linker is highlighted in green.

reproduce this human-designed approach by providing the same five fragments as input and tasking it with generating novel molecular combinations. Remarkably, our model successfully generated four molecules among the twenty four molecules present in the submission series, as shown in figure 6. Our model can efficiently identify a variety of drug-related functional groups, including pyridine rings, amide bonds, and others. This achievement demonstrates our model’s sophisticated ability to identify mergeable sites and recognize the structural compatibility between different fragment scaffolds.

The success in reproducing four out of five fragments highlights several key strengths of our approach: (1) the model’s capacity to recognize pharmacophoric features that are essential for maintaining binding affinity, (2) its ability to identify chemically feasible linking strategies between fragments. This performance validates our model’s potential as a valuable tool for fragment-based drug design, particularly in its ability to automate and accelerate the traditionally labor-intensive process of fragment merging and optimization.

A.2 Pre-training and finetuning dataset

For pretraining, we utilized the ZINC dataset, selecting only molecules categorized as Standard, In-Stock, and Drug-Like, which resulted in approximately 11 million compounds. Preprocessing involved two straightforward steps: (i) canonicalizing SMILES strings using `Chem.MolToSmiles(Chem.MolFromSmiles(mol), True)`, and (ii) discarding molecules with empty scaffold SMILES. These procedures are also incorporated into the HuggingFace data repository.

For finetuning, we employed 1 million compounds from the ZINC15 dataset, each docked to the RTCB protein (PDB ID: 4DWQ) implicated in human cancer, as provided by the Cancer dataset of Liu et al. [23], and applied them consistently across all baselines. In each finetuning epoch, a random subset of the desired number of molecules was sampled.

For evaluation and experimentation, we used 1,280 molecules from the Cancer dataset that did not overlap with the finetuning set.

A.3 Pretraining

Incremental pre-training. In the first phase, we train a GPT model exclusively on molecular data using Causal Language Modeling (CLM). In this autoregressive setup, the model learns to predict the




Figure 6: Four merging cases were generated, with all generated molecules contained within the 24 molecules of submission TRY-UNI-714a760b. Different fragments are marked in red and blue, while the merged overlapping regions are highlighted in green.

next token in a sequence conditioned only on the preceding tokens. The Phase 1 training corpus is constructed as follows:

$$C_{\text{Phase 1}} = \left\{ [BOS], \langle L \rangle, \underbrace{y_1, \dots, y_T}_{\text{target molecule Y}}, [EOS] \right\}, \langle L \rangle \text{ is the token for ligand.} \quad (7)$$

In the second phase, we build upon the GPT model developed in Phase 1, which has already demonstrated strong performance in molecular generation. Here, the training objective is extended to focus on fragment–molecule relationships, where the model is trained on pairs of fragments and their corresponding complete molecules using Causal Language Modeling (CLM). The Phase 2 training corpus is constructed as follows:

$$C_{\text{Phase 2}} = \left\{ [BOS], \langle p_1 \rangle, \underbrace{f_{11}, \dots, f_{1T}}_{\text{fragment } F_1}, \langle p_2 \rangle, \underbrace{f_{21}, \dots, f_{2T}}_{\text{fragment } F_2}, \langle L \rangle, \underbrace{y_1, \dots, y_T}_{\text{target molecule Y}}, [EOS] \right\}, \quad (8)$$

where $\langle F_1, F_2 \rangle$: fragments. Consequently, the model is able to generate the corresponding complete molecule when provided with a pair of fragments as the input prompt.

A.4 Surrogate model

The surrogate model [33] adopts a simplified BERT-style transformer architecture, widely used in natural language processing. In this framework, tokenized SMILES strings are embedded with positional encodings and passed through a stack of five transformer blocks. Each block consists of a multi-head attention mechanism with 21 heads, followed by dropout, layer normalization with residual connections, and a feedforward subnetwork. The feedforward subnetwork is composed of two dense layers, again followed by dropout and layer normalization with residual connections. After the transformer stack, a final feedforward layer produces the predicted docking score. The validation performance achieved r^2 values of 0.73 on the RTCB dataset.

A.5 Computing infrastructure and wall-time comparison

We trained our docking surrogate models using 4 nodes where each node contains CPUs (64 cores) and 4 A100 GPU nodes. The training time for each model was approximately 3 hours. We conducted other RL experiments on a cluster that includes CPU nodes (approximately 280 cores) and GPU nodes (approximately 110 Nvidia GPUs, ranging from Titan X to A6000, set up mostly in 4- and 8-GPU configurations). Based on the computing infrastructure, we obtained the wall-time comparison in Table 2 as follows.

Methods	Total Run Time
Pretraining	36h
Finetuning	30h

Table 2: Wall-time comparison between different methods.

A.6 Hyperparameters and architectures

Table 3 provides a list of hyperparameter settings we used for our experiments.

For experimentation, we use a set of 1,280 molecules from the RTCB dataset, selecting those with docking scores in the range of $[-14, -6]$, following the range reported in [40].

When computing the average normalized reward of the original molecules—without incorporating similarity constraints—we assign equal weights of 0.25 to each of the four objectives: docking, drug-likeness, synthesizability, and solubility.

Parameter	Value
Pretraining	
Learning rate	$5 \times e^{-5}$
Batch size	24
Optimizer	Adam
# of Epochs for Training First Phase	10
# of Epochs for Training Second Phase	10
Model # of Params	124M
Finetuning	
# of molecules generated K	4000
# of Epochs N	50
Batch Size	16
Learning rate	$5 \times e^{-5}$
Generation	
TopK	[10, 15, 20]
TopP	[0.85, 0.9, 0.95]
Shared	
Tamimoto Similarity Fingerprint Size	1024

Table 3: Hyperparameters for pretraining and generation.