De Novo Antibiotic-Like Molecule Design via Descriptor-Guided Patch-Based GANs

Syed Rameez Naqvi, Md Mostafizer Rahman and Lu Peng

Department of Computer Science, Tulane University
6823 Saint Charles Ave, New Orleans, LA 70118 USA
{snaqvi, mrahman9, lpeng3}@tulane.edu

Abstract

The rise in antimicrobial resistance necessitates rapid discovery of novel antibiotics, but the majority of generative pipelines are unable to produce clinically viable candidates. The structural and topological complexity of actual antibiotics is not captured by current models, which lack pharmacological depth and are frequently trained on datasets such as QM9 and ZINC. We present a modular, descriptor-guided framework for antibiotic-like molecule design that combines a property-aligned β-VAE for interpretable encoding, a descriptor-to-latent conditioner for controllable sampling, and a patch-based graph generator for fragment-wise synthesis. Trained on a curated subset of ChEMBL containing clinically validated antibiotics, our framework supports end-to-end generation from RDKit descriptors to final molecules, with adversarial refinement for topological realism. Beyond favorable ADMET profiles, our method establishes *joint satisfaction reporting* of drug-likeness thresholds as a benchmark standard, resists mode collapse under 50k-sample stress tests, and surfaces ligands that outperform ciprofloxacin and co-crystal references in docking assays. These results highlight chemically meaningful, pharmacologically informed generation – overcoming limitations of black-box pipelines and general-purpose datasets.

Introduction

Antimicrobial resistance continues to pose threat to global health by decreasing the effectiveness of numerous antibiotics (Ferrara et al. 2024). The development of new antibiotics has slowed considerably in spite of this urgency, in part due to the high failure rate and high cost of drug development. Only a small fraction of candidates ever reach patients, and development timelines often exceed a decade (WHO 2022). Consequently, AI-driven approaches for accelerating pharmacologically viable antibiotic discovery have gained increasing attention (Swanson et al. 2024).

Generative Adversarial Networks (GANs) are established tools for exploring large chemical spaces and generating molecules with desired properties for drug design (Tong et al. 2021). Most generative models in use, however, are evaluated using datasets such as QM9 (Ramakrishnan et al. 2014) and ZINC (Tingle et al. 2023), which, although computationally efficient, do not provide the pharmacological depth required for practical antibiotic discovery (Cauchy, Leguy, and Da Mota 2023). ZINC concentrates on docking-ready compounds without experimental target or mechanism data, whereas QM9 consists of small, synthetically enumerated molecules without bioactivity annotations (Glavatskikh et al. 2019). In contrast, ChEMBL offers a manually curated repository of drug-like molecules with extensive annotations including IC_{50}, Ki, mechanisms of action, and therapeutic targets (Zdrazil et al. 2023). We focus on the generation of antibiotic-like molecules using a curated subset of ChEMBL consisting of clinically validated antibiotics. Throughout, *antibiotic-like* refers to molecules within this chemical domain; all validation remains in silico and does not constitute experimental evidence of biological activity.

To this end, we propose a property-aligned generative framework that conditions latent molecule representations on RDKit descriptors (Landrum 2013), including quantitative estimate of drug-likeness (QED), the octanol-water partition coefficient (logP), and the synthetic accessibility score (SA). These descriptors are mapped to latent vectors and used to generate fixed-size graph patches – interpretable molecular fragments – that are subsequently assembled into full molecules. Unlike prior GAN pipelines (De Cao and Kipf 2018; Guimaraes et al. 2017; Wei et al. 2023), which optimize or report properties independently, we establish *joint satisfaction of pharmacologically meaningful thresholds* as a benchmark. We report the fraction of molecules simultaneously meeting QED > 0.6, SA < 5, and logP $\in [-0.5, 5.0]$, a population-level metric absent from existing baselines. Our descriptor-conditioned patch generator not only provides interpretable fragment-level control but also mitigates the mode collapse common in atom-level GANs (Tang et al. 2024). By this, we introduce *multi-property alignment* as an evaluation metric, which helps yield hundreds of topologically diverse, jointly optimized antibiotic-like molecules with favorable ADMET profiles (Daoud et al. 2021).

Unlike purely end-to-end generative models, our framework integrates structured chemical knowledge at multiple stages. Descriptor-level priors (QED, logP, SA) impose symbolic constraints aligned with medicinal chemistry practice; RDKit-based valence rules ensure graph validity; and patch aggregation encodes domain-informed structural composition rather than unconstrained full-graph sampling. This

modular design combines learned components with rule-based chemical constraints, embedding domain expertise directly into the generative process rather than relying solely on emergent behavior from data. Our main contributions are:

- A curated dataset of 4,607 antibiotic-like molecules from ChEMBL, annotated with drug-likeness properties and RDKit descriptors.
- A modular framework combining a β-VAE, descriptor-to-latent conditioner, and patch-based generator for interpretable, property-guided molecule design.
- Establishing *joint satisfaction reporting* (QED > 0.6, SA < 5, logP $\in [-0.5, 5.0]$) as a practical benchmark for multi-property drug design, exposing gaps in prior baselines.
- A 50k-sample stress test demonstrating 72% unique scaffolds and robustness against mode collapse.
- In silico validation via QSAR screening and docking against antibacterial targets (DHFR, DNA gyrase), identifying candidates outperforming reference ligands.

Background & Related Work

Fragment- & Graph-Based Models: SMILES-based sequence models, such as RNNs and Transformers, suffered from syntactic invalidity and lacked structural interpretability (Andronov et al. 2025). While graph-based models, such as MolGAN (De Cao and Kipf 2018) and GraphAF (Shi et al. 2020), demonstrated improved chemical validity via direct molecular graph generation, most of them suffered mode collapse and failed to provide fine-grained control over structural components. Recent diffusion-based models such as EDM (Hoogeboom et al. 2022) showed strong performance in generating 3D molecular structures with high stability. While powerful, these models were primarily unconditioned and focused on geometric fidelity rather than pharmacological control. They lacked mechanisms for descriptor-level property steering and did not operate in a fragment- or scaffold-based regime. To introduce modularity, hierarchical models such as Junction Tree (JT)- VAE (Jin, Barzilay, and Jaakkola 2020) and DeepMGM (Bian and Xie 2022) decomposed molecules into scaffold–fragment pairs, enabling improved scaffold retention. These models, however, often relied on fixed fragment vocabularies. ScaffoldGVAE (Hu et al. 2023) introduced multi-view graph encoding to enhance scaffold-level diversity and pharmacophoric consistency.

Domain-Specific Molecular Design: Most existing generative models applied rewards post hoc – MolCycleGAN (Maziarka et al. 2020) and InstGAN (Tang et al. 2024) optimized generated molecules towards target properties using reinforcement learning or adversarial training. However, they lacked explicit alignment between input descriptors and molecular representations. This shortcoming has made in-generation property control a key challenge in molecule generation. While most existing models were benchmarked on general-purpose datasets such as ZINC and QM9 (Glavatskikh et al. 2019), recent efforts explored target-specific molecule generation. For example, DeepMGM (Bian and Xie 2022) focused on CB2-targeted libraries, and RuSH (Rossen et al. 2024) introduced a reinforcement learning strategy for scaffold hopping based on pharmacophoric similarity.

Despite recent advances, key limitations persist in the molecular generation literature. First, most models prioritize structural validity and property optimization but lack mechanisms for interpretable, descriptor-guided control during generation (Du et al. 2024). Second, fragment-based methods often rely on discrete vocabularies or scaffold libraries, limiting generative flexibility (Voloboev 2024). Third, while some models incorporate property rewards, few align molecular latent spaces explicitly with interpretable chemical descriptors (Haddad et al. 2025). Finally, not many works focus on antibiotic discovery, where new chemical diversity is urgently needed (Schuh, Hesse, and Sieber 2025; Chen, Yu, and Gao 2023). This motivates the development of a modular, property-aligned framework that supports controllable, fragment-level generation tailored to a clinically relevant therapeutic space.

Proposed Framework

Preliminaries

Inspiration and Adaptation: Our design is inspired by modular, property-aware molecule generation concept. Fragment-based models such as JT-VAE (Jin, Barzilay, and Jaakkola 2020) and DeepMGM (Bian and Xie 2022) motivate our patch-level synthesis approach; we adopt a continuous latent representation instead of fixed vocabularies though (Fig. 1(a)). Inspired by Mol-CycleGAN (Maziarka et al. 2020) and InstGAN (Tang et al. 2024), we employ β-VAE (Higgins et al. 2017) to align learned latents with RDKit descriptors for controllable generation. Finally, ScaffoldGVAE (Hu et al. 2023) inspires scaffold diversity and coherent graph reconstruction in our framework.

Comparison with Related Works: Table 1 summarizes key attributes for various molecule generation baselines. While prior methods offer strong validity or partial property control, few support fragment-based decoding, latent interpretability, or clinically relevant descriptor conditioning. In contrast, our framework integrates these capabilities into a unified pipeline aligned with ADMET evaluation and scaffold-aware synthesis.

Dataset Construction and Feature Extraction: We curate a representative dataset of antibiotic molecules from ChEMBL (Zdrazil et al. 2023) by filtering drug indications for the keyword "antibiotic". Molecules with invalid SMILES or fewer than 10 heavy atoms are removed, yielding a high-quality set of 4,613 compounds. For each molecule, we compute three core properties: QED, logP and SA. These properties are calculated using RDKit and a publicly available SA scoring module. Each molecule m is represented by a 3-dimensional feature vector:

$$\mathbf{x}_m = [\text{QED}(m),\ \text{logP}(m),\ \text{SA}(m)]$$

We then enrich this filtered set with 200+ RDKit descriptors (Landrum 2013) including physicochemical and topological features. Let $\mathbf{d}_m \in \mathbb{R}^D$ denote the RDKit descriptor

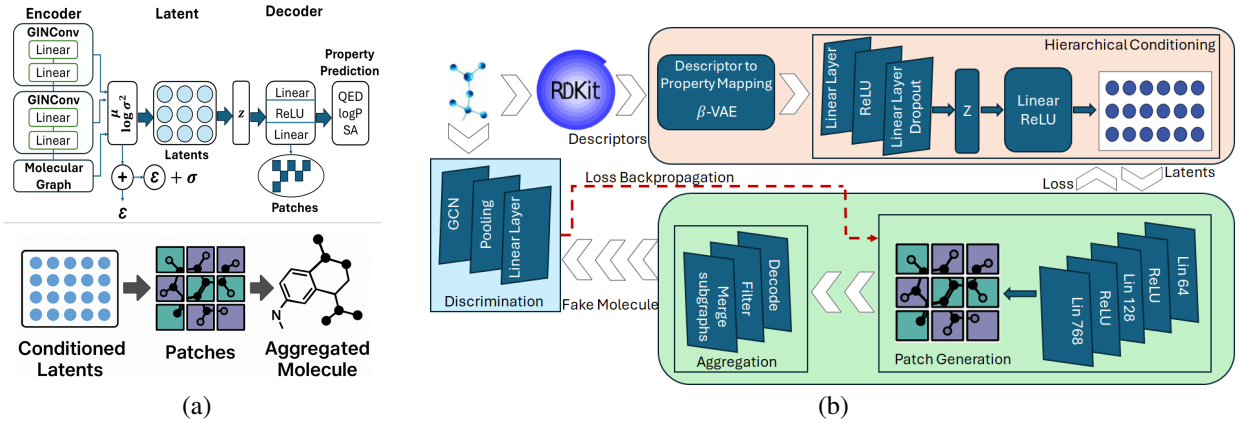

Figure 1: (a) Proposed molecule construction, (b) End-to-end framework overview

Model	Dataset	Validity	Property Control	Latent Interpretability	Patch-Based Generation
MolGAN*	ZINC	∼98%	✗	✗	✗
GraphAF†	QM9	100%	✓(flow)	✗	✗
DeepMGM‡	ZINC	100%	✓	✓(vocab)	✗
EDM§	QM9	100%	✗	✗	✗
RuSH¶	ZINC	∼99%	✓(target-based)	✗	✗
HierVAEζ	Polymers	100%	✗	✗	✓(Motif)
Proposed	ChEMBL	100%	✓(desc-cond)	✓(β-VAE)	✓

Table 1: Comparison of molecular generation methods across key modeling properties
Legend: *(De Cao and Kipf 2018), †(Shi et al. 2020), ‡(Bian and Xie 2022), §(Hoogeboom et al. 2022), ¶(Rossen et al. 2024), ζ(Jin, Barzilay, and Jaakkola 2020)

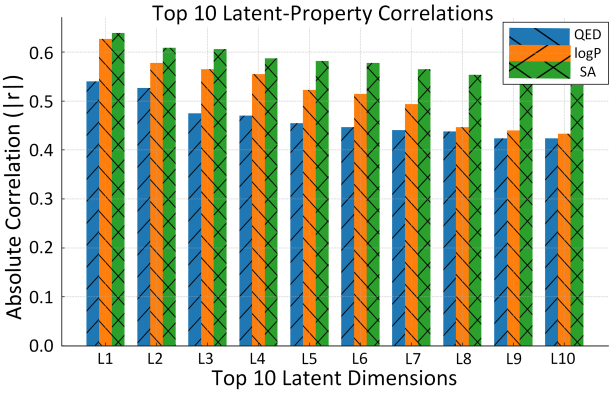

Figure 2: Latent dim. & mol. props. correlation

vector and $\mathbf{p}_m \in \mathbb{R}^3$ the property vector, then the final descriptor matrix is:

$$\mathbf{x}_m = [\mathbf{d}_m \parallel \mathbf{p}_m] \in \mathbb{R}^{D+3}$$

Although trained on antibiotics, our framework generalizes to any therapeutic class for which meaningful descriptors are available.

Pipeline Components

The lower portion of Fig. 1(a) conceptually illustrates the generative flow: conditioned latent dimensions are first translated into fixed-size graph patches, which represent localized structural motifs. These patches are then merged and decoded into full molecular graphs under valence and connectivity constraints. Fig. 1(b) presents an end-to-end framework showing how RDKit descriptors are mapped to latent property vectors via a β-VAE and Conditioner. These guide patch generation and aggregation into molecules, later evaluated by a graph discriminator. Adversarial feedback refines patch synthesis.

Variational Autoencoder for Latent–Property Modeling: To model molecular representations with interpretable, property-aligned latents, we train a β-VAE (Higgins et al. 2017) with a supervised head for QED, logP, and SA prediction. The encoder is a 2-layer GIN (Xu et al. 2019), which projects input graphs to latent parameters $\mu, \log \sigma^2$, for latent vector sampling:

$$z \sim \mathcal{N}(\mu, \mathrm{diag}(\sigma^2)), \quad z = \mu + \sigma \odot \epsilon$$

The decoder (MLP) reconstructs molecular patches, while an auxiliary head regresses property values. Total loss is:

$$\mathcal{L} = \beta \cdot \mathcal{L}_{\mathrm{KL}} + \lambda \cdot \mathcal{L}_{\mathrm{prop}}, \quad \mathcal{L}_{\mathrm{prop}} = \sum \mathrm{MSE}_{\mathrm{QED,\ logP,\ SA}}$$

with $\beta = 0.1$, $\lambda = 5.0$. Pearson correlations between latent dimensions and molecular properties (Fig. 2) reveal consistent, interpretable alignment trends. We extract the top-k most correlated latents per property to support targeted conditioning and traversal. To ensure alignment between VAE latents and descriptors, we intersect molecules present in both representations. Invalid features or those with >99% constant values are removed, yielding a clean matrix $X \in \mathbb{R}^{n \times d}$ – used for downstream analyses:

$$X = \left[x_1^\top \ \cdots \ x_n^\top\right]^\top, \quad n = 4607,\ d = 188$$

Learning Descriptor-to-Latent Mappings via Conditioning: To enable property-driven molecule generation, we train a *Conditioner* – a 3-layer MLP mapping RDKit descriptors $\mathbf{x} \in \mathbb{R}^d$ to a 9-dimensional latent vector $\mathbf{z}_{\text{top}} \in \mathbb{R}^9$ aligned with molecular properties. These 9 latent dimensions are selected based on the highest Pearson correlations with QED, logP, and SA.

$$f_{\text{cond}} : \mathbf{x} \mapsto \mathbf{z}_{\text{top}}, \quad \mathbf{z}_{\text{top}} = W_3 \cdot \text{ReLU}(W_2 \cdot \text{ReLU}(W_1 \cdot \mathbf{x}))$$

We train a descriptor-to-latent conditioner to map physicochemical descriptors into the learned latent space, enabling molecule generation from descriptor inputs without requiring the original molecular graph. In practice, the trained conditioner exhibits strong alignment between predicted and target latent representations, supporting its role as a bridge between descriptor space and property-aligned latent embeddings. The β-VAE and conditioner architectures follow standard design choices and are illustrated in Fig. 1(a-top).

Empirically, multiple latent dimensions exhibit consistent monotonic relationships with QED, logP, and SA, enabling targeted traversal and descriptor-aligned generation without post hoc reward shaping.

Patch Generator via Latent Supervision: We train a patch generator $\mathcal{G}$ to synthesize node-feature matrices from descriptor-derived latent vectors. The generator is a 3-layer MLP that receives a 9-dimensional vector $\mathbf{z}_{\text{top}}$ from Conditioner, and outputs a latent graph patch $\hat{\mathbf{P}} \in \mathbb{R}^{48 \times 16}$. To supervise training, we pass $\mathbf{z}_{\text{top}}$ through a fixed decoder from the pretrained VAE to obtain a reference patch. The generator minimizes the MSE loss:

$$\mathcal{L}_{\text{gen}} = \left\| \mathcal{G}(\mathbf{z}_{\text{top}}) - \hat{\mathbf{P}} \right\|_2^2$$

We train for 25 epochs with Adam (lr $= 10^{-3}$), saving the best model based on validation loss. This supervised generator bridges tabular descriptors and molecular substructures, enabling controllable synthesis of latent graph embeddings.

Patch Aggregation and Chemical Constraints. Generated latent patches are decoded into full molecular graphs using a constraint-aware aggregation procedure designed to preserve chemical validity without relying on a fixed fragment vocabulary. Aggregation enforces atom valence limits, degree constraints, and ring consistency, while favoring the formation of chemically plausible scaffolds through local connectivity patterns. By operating at the patch level rather than atom-by-atom generation, this strategy reduces invalid intermediate states and mitigates mode collapse commonly observed in graph-based GANs, while maintaining flexibility to assemble diverse molecular topologies.

Graph-Based Discriminator for Realism Detection: To distinguish ChEMBL molecules from generated ones, we train a graph-based binary discriminator using a two-layer graph convolutional network (GCN) followed by global mean pooling and fully connected layers in a standalone warm-up phase with equal branches of the two types of graphs. Atom-level features and edge indices are derived from SMILES strings using RDKit and PyTorch Geometric. Given input $\mathbf{x}$ and edge set $\mathcal{E}$:

$$\mathbf{h}_1 = \text{ReLU}(\text{GCNConv}_1(\mathbf{x}, \mathcal{E}))$$

$$\mathbf{h}_2 = \text{ReLU}(\text{GCNConv}_2(\mathbf{h}_1, \mathcal{E}))$$

The pooled embedding is passed through two linear layers to produce a logit, optimized with binary cross-entropy:

$$\mathcal{L}_{\text{disc}} = \text{BCEWithLogitsLoss}(\hat{y}, y)$$

The pretrained discriminator serves both as a standalone molecule filter and as a feedback module in adversarial training. Intermediate graph embeddings are also used for diversity analysis.

Adversarial Fine-Tuning through GAN Training: We employ GAN-based training to improve patch generation, where the generator minimizes binary cross-entropy loss to learn to produce molecules that are indistinguishable from the real samples – as determined by the discriminator:

$$\mathcal{L}_{\text{gen}} = \text{BCE}(D(G(z)), 1)$$

where latent patches z are aggregated into molecules and scored by the discriminator during each iteration, which alternates between discriminator- and generator-updates. This loop maintains property alignment while enhancing topological realism. Algorithm 1 describes the framework's operation in a four-phase VAE-GAN pipeline.

Algorithm 1: Descriptor-Conditioned Patch-Based Molecule Generation

Require: Antibiotic-like molecular dataset, descriptor vectors $\mathbf{x}$
Ensure: Generated molecules with controlled physicochemical properties

1: **Phase I (Latent Pretraining):** Train a β-VAE encoder–decoder to map molecular descriptors to a smooth latent space and reconstruct fixed-size molecular patches.
2: **Phase II (Discriminator Warmup):** Train a graph discriminator to distinguish real molecules from aggregated patch-based generations.
3: **Phase III (Patch Generation):** Learn a descriptor-conditioned generator that maps $\mathbf{x}$ to latent patches using adversarial feedback from the fixed discriminator.
4: **Phase IV (Assembly & Evaluation):** Decode generated patches into full molecular graphs using a valence- and ring-aware aggregator, then evaluate validity and physicochemical properties (QED, SA, logP).

Results & Discussions

All experiments use a fixed random seed for reproducibility. We use a curated subset of 4,607 ChEMBL antibiotic SMILES as input space, generating ∼9,000 molecules per experiment. Chemical validity, uniqueness, and novelty are assessed with RDKit, alongside physicochemical properties (QED, logP, SA). Beyond these core metrics, we evaluate ADMET profiles (SwissADME/pkCSM heuristics), perform large-scale robustness tests, and conduct biological validation via docking and QSAR screening.

Diversity and Latent Structure Diagnostics: To assess whether the generator's output diversity reflects meaningful conditioning (not random dispersion), we compare three conditions: (1) untrained generator initialized with random weights, (2) trained generator receiving random latent vectors (with the conditioner removed), and (3) the full pipeline with descriptor-conditioned latent input. For each case, we visualize the t-SNE embeddings of the generated patches, see Fig. 3. The untrained generator produces patches that are highly diverse (cosine similarity spanning 0.3–0.9), but the diversity is unstructured, as shown in the t-SNE scatter. This confirms that architectural priors alone are insufficient for semantically meaningful generation. In the ablation condition (trained generator with random latent vectors), patches remain diverse, but no consistent gradient is observed with respect to QED, logP, or SA. Only in the fully conditioned case do we observe consistent improvements in downstream multi-property satisfaction; the t-SNE plots are included as a qualitative diagnostic of latent structure and diversity rather than as definitive evidence of property gradients. Topological diversity and complexity metrics further confirm the generator's ability to produce pharmacologically realistic compounds. Note that we do not retrain the full GAN pipeline without the conditioner, since the ablation of property alignment is most directly attributable to the latent-level input. Adding GAN dynamics will introduce orthogonal sources of instability and obscure the conditioner's contribution.

GAN Checkpoint: To identify the optimal checkpoint, we analyze property scores for molecules generated per epoch, and develop a custom reward function to rank epochs by molecular quality:

$$\text{Reward} = 2 \times \text{QED} - 0.6 \times \left(\frac{\max(\text{SA} - 3,\ 0)}{4} \right) - 0.4 \times \left(\frac{\max(|\text{logP} - 2.5| - 1,\ 0)}{2} \right)$$

where SA and logP penalties are softly applied outside favorable ranges (SA $>$ 3.0, logP far from 2.5), and all values are normalized. Coefficients reflect pharmacological relevance and empirical property distributions: QED has the tightest dynamic range, while SA and logP penalties are variance-scaled to avoid dominating the score. A 20% sensitivity analysis altered epoch rankings by less than 2%, confirming robustness. The epoch with the highest mean reward is selected for all downstream analyses.

Molecule Generation from RDKit Descriptors: To evaluate generative capacity, we perform large-scale inference by sampling molecules from RDKit descriptors. Using trained conditioner and generator checkpoints, each descriptor is projected into a 9-dimensional latent space, from which molecular patches are generated, which are decoded into candidate molecules, ensuring valence-aware graph reconstruction. For each descriptor, two to four molecules are generated (Mols/Desc), subject to a global cap of 16,500. Molecules are then validated and scored using standard cheminformatic metrics (QED, logP, SA). A total of 16,268 valid molecules are obtained, forming the basis for subsequent assessments. Average inference time per molecule and total trainable parameters for various classical and proposed methods are shown in Tab. 2. Our method achieves faster inference than most baselines, despite modular property control and fragment-level synthesis.

Evaluation of Generated Molecules: After validating SMILES strings, we compute QED, logP, and SA, along with uniqueness, novelty, and Tanimoto-based topological diversity. A composite *goodness* score penalizes molecules with SA $>$ 4 or logP $>$ 4, providing a scalar ranking:

$$\text{Goodness} = \text{QED} - 0.2 \times \text{SA}_{\text{penalty}} - 0.1 \times \text{logP}_{\text{penalty}}$$

where penalties are linearly scaled above thresholds. A t-SNE projection confirmed that generated molecules align well with dataset's distribution. Next, to evaluate the pharmacokinetic and safety profiles of the generated molecules, we implement an ADMET scoring script inspired by SwissADME and pkCSM heuristics. The resulting profiles are favorable: molecular weights cluster around 250–325 Da, logP values between 2.5–5, and TPSA within 20–60 A^2, supporting bioavailability. Most molecules pass key filters and are predicted to be non-toxic. Notably, Composite_ADMET scores peak above 0.9, confirming the generated compounds' strong pharmacological potential.

Method	Model	Gen.	Param.	T (s)
MLP-VAE†	VAE	One Shot	360,448	0.04
E-NFs‡	Flow	One Shot	647,117	0.27
G-ScheNet*	Sampling	Auto Reg	902,111	0.41
G-SphereNet$'$	Flow	Auto Reg	3,148,095	0.55
EDM$^{\circ}$	Diffusion	One Shot	5,340,921	0.86
JT-VAE^{+}	VAE	Two-Stage	5,700,000	0.12
GraphAF^{-}	Flow	Auto Reg	2,100,000	0.18
MolGAN$^{\#}$	GAN	One Shot	1,200,000	0.09
L-MolGAN§	GAN	One Shot	1,350,000	0.11
GCPN¶	RL	Auto Reg	1,800,000	0.25
Proposed	GAN	One Shot	462,089	0.06

Legend: † (Kingma, Welling et al. 2013), ‡ (Garcia Satorras et al. 2021), * (Luo and Ji 2022), $'$ (Hoogeboom et al. 2022), $^{\circ}$ (Gebauer, Gastegger, and Schütt 2019), $^{+}$ (Jin, Barzilay, and Jaakkola 2020), $^{-}$ (Shi et al. 2020), $^{\#}$ (De Cao and Kipf 2018), § (Tsujimoto et al. 2021), ¶ (You et al. 2018)

Table 2: Parameter count and inference time per mol

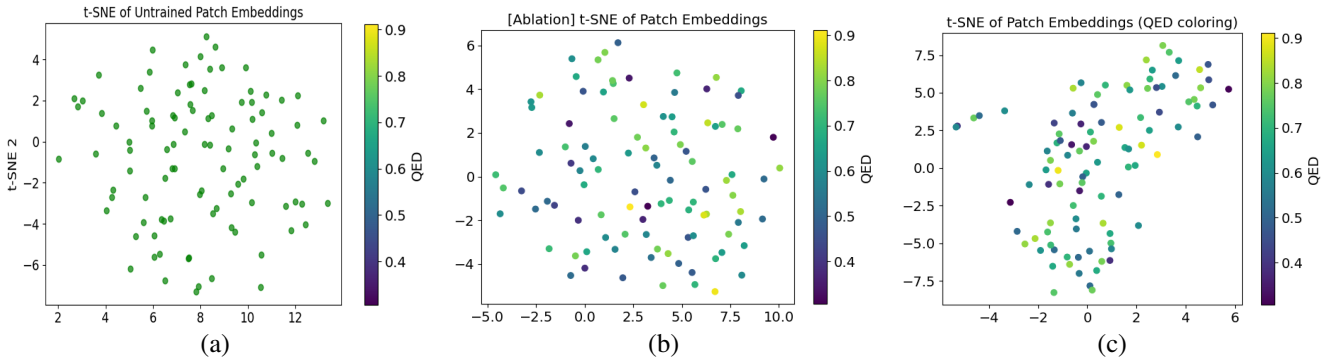

Figure 3: t-SNE embeddings of generated patch representations colored by QED score. (a) untrained generator (random weights), (b) trained generator with random latent input (conditioner removed), and (c) full pipeline with descriptor-conditioned latent input. We use t-SNE as a qualitative diagnostic of latent structure and diversity; property satisfaction is quantified separately via thresholded and joint metrics in Table 3

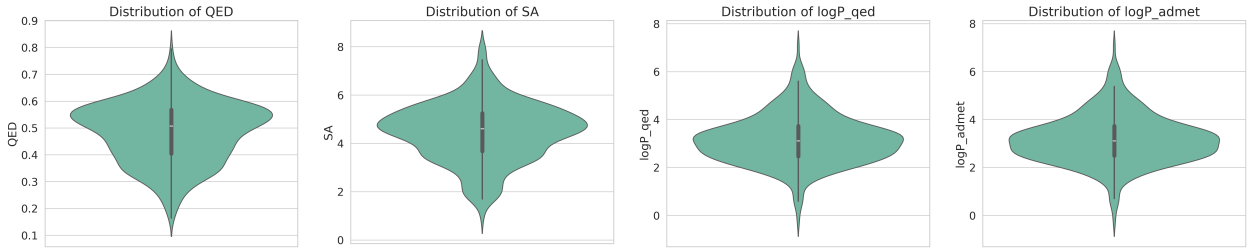

Figure 4: Distribution of key molecular properties among the generated molecules. QED, SA score, logP (from generation), and logP (recomputed during ADMET profiling)

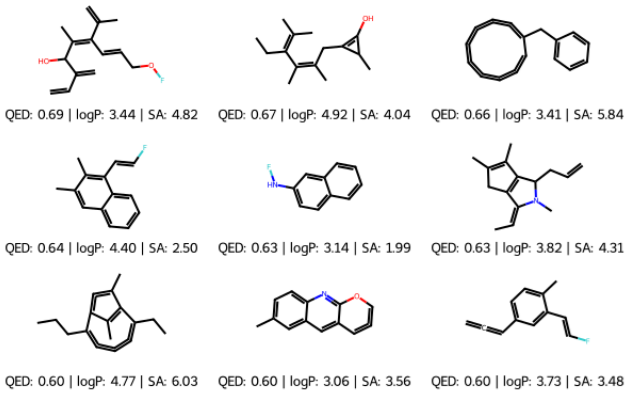

Figure 5: A randomly selected grid of generated molecules

For generated molecules' pharmacological relevance, we visualize the distribution of key molecular properties using violin plots (Fig. 4. While the QED values – clustered around 0.80–0.85 – indicate strong drug-likeness, the SA scores exhibit a broad distribution (from 3 to 8), reflecting a balance between synthetic tractability and topological novelty. The agreement between generator logP and recomputed logP values confirms internal consistency, while the TPSA and logS distributions are centered in bioavailable regions. Most notably, the composite ADMET scores show a peak above 0.85, supporting the molecules' overall chemical viability. The generated samples, Fig. 5, beside endorsing QED and SA observed through violin plots, showcase topologically diverse molecules. Note that most logP values fall between 2.8 and 4.9 – usually considered acceptable for oral bioavailability. The observed ring systems, branching patterns, and functional groups, such as amines, alcohols, ethers, and halogens, confirm high topological diversity. We further calculate Frechet ChemNet Distance (FCD) (Preuer et al. 2018) using ECFP4 fingerprints to evaluate the similarity between generated and reference molecules. Multivariate Gaussians are used to approximate 2048-bit features from 4,607 generated and 5,000 ChEMBL compounds, which are then reduced to 100 principal components. An FCD of 0.1302 indicates strong alignment between chemical spaces.

Property Traversal & Pareto Analysis: We conduct a property traversal experiment focusing on QED (↑), SA (↓), and logP ($\in [-0.5, 5.0]$) by projecting generated molecules into the (QED, logP) plane with SA as a color scale. Pareto fronts (Fig. 6) reveal trade-offs: unconstrained (black) vs. chemically realistic (red, SA $\leq$ 6.0, logP bounded). This highlights how hard constraints filter impractical candidates confirming that our patch-based conditioning enables smooth property control and Pareto-optimal exploration, which classical GANs/VAE models fail to expose.

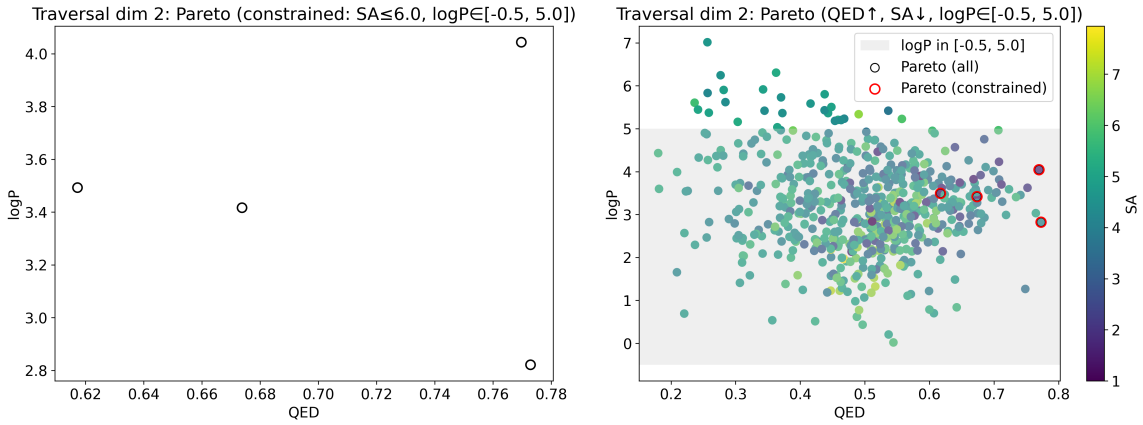

Figure 6: Property traversal experiment on latent dimension 2. Left: Pareto front only (constrained). Right: Full scatter plot with unconstrained Pareto front (black) and constrained Pareto front (red).

Mode Collapse Stress Test: To test robustness against mode collapse, we generate 50,000 molecules under identical latent seeds and descriptor-conditioned settings. Despite this extreme sampling, the model sustained diversity: all molecules are chemically valid, no duplicates are detected, and ∼72% exhibit distinct Bemis–Murcko scaffolds (36k unique scaffolds). Internal similarity is low (mean Tanimoto ∼0.14), confirming the generator does not degenerate into repetitive outputs. Property distributions (Fig. 7) further show broad coverage across QED and logP ranges, supporting chemical realism alongside topological diversity. Although SA is not computed in this experiment due to runtime constraints, the QED–logP scatter plot (Fig. 7c) further illustrate a wide coverage of physiochemical space, with no collapse into narrow property bands.

Bio-validation: Beside establishing topological diversity and drug-likeness, practical relevance of molecular generative frameworks also requires evidence of biological activity. We conduct preliminary in-silico bio-validation through docking (AutoDock Vina) and QSAR screening. Docking evaluates whether generated compounds can plausibly bind known antibacterial targets, whereas QSAR models provide a statistical proxy for activity across diverse scaffolds.

On *E. coli* DHFR (6XG5), 281 ligands are docked, with the best ligand scoring –8.41 kcal/mol, outperforming the redocked co-crystal (≈–7.70 kcal/mol). On *S. aureus* DNA gyrase (2XCT), 9 ligands surpass ciprofloxacin (co-crystal ≈–9.50 kcal/mol), with the best achieving –12.05 kcal/mol. These results demonstrate that our descriptor-to-latent patch generator enriches high-affinity candidates across two antibacterial targets. A lightweight QSAR model further flags 12 generated molecules as both drug-like and biologically plausible. Fig. 8 summarizes these results,which establish our framework as a practical step toward AI-driven antibiotic discovery by achieving not only chemical realism but also biological plausibility.

Comparison with Baseline Models: We evaluate post-GAN generation quality in comparison to recent baselines that report similar metrics, Table 3. While many published models claim to generate valid or drug-like molecules, only InstGAN (Tang et al. 2024) and L-MolGAN (Tsujimoto et al. 2021) report thresholded statistics or sufficient breakdowns to allow a fair comparison. Moreover, other diffusion/flow models like GraphAF (Shi et al. 2020) or EDM (Hoogeboom et al. 2022) focus on autoregressive/full-graph generation and do not support descriptor conditioning, making direct comparison non-trivial. We include them for context only, as their objectives and conditioning mechanisms differ fundamentally. While they achieve high validity, they do not report joint satisfaction nor descriptor-level conditioning; thus our benchmark exposes a gap.

InstGAN reports QED = 0.93, SA = 0.99, logP = 0.89, but these apply only to the Top-1000 molecules and are reported independently for each property. For a leveled comparison, we additionally report statistics for our own Top-1000 molecules ranked by the composite goodness score. In this subset, 100% satisfy QED > 0.6, 92.4% satisfy SA < 5, and 99.7% satisfy logP $\in [-0.5, 5]$, with 92.1% MPS. This Top-1000 analysis mirrors the ranking-based reporting strategy used by InstGAN while preserving our multi-property evaluation framework.

Similarly, L-MolGAN provides single-point values under various property-weighting schemes (e.g., $\lambda = 0.1$). However, none reports how many molecules meet multi-property thresholds. In contrast, beyond Top-1000 reporting, our method provides granular population-level filtering, reporting detailed distributions and absolute counts of molecules that meet pharmacologically meaningful cutoffs. We generate 1821 such molecules at the 4 molecules per descriptor (MPD) level, offering a practical measure of promising candidates. Although aggregate QED and SA rates remain conservative, explicit descriptor control combined with latent conditioning enables the generation of structurally diverse molecules that jointly satisfy multiple physicochemical constraints. Joint satisfaction reporting is adopted to reflect practical drug discovery settings, where candidate molecules must simultaneously meet several thresholds rather than op-

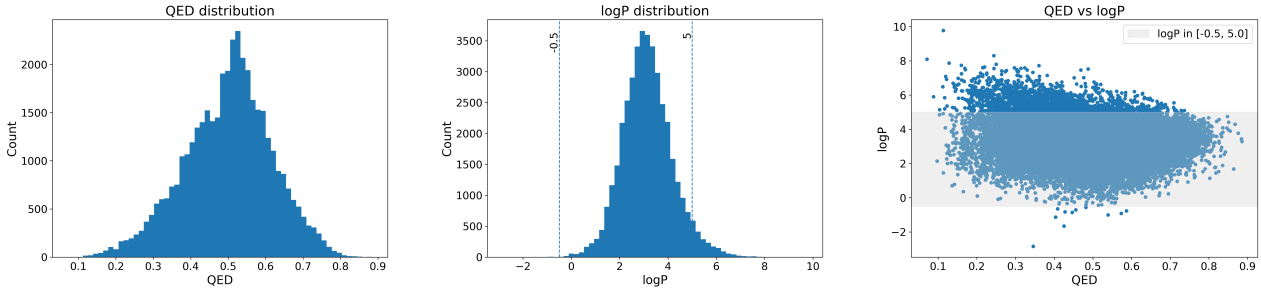

(a) (b) (c)

Figure 7: Mode collapse stress test: (a) Distribution of QED values, (b) distribution of logP values, and (c) joint QED–logP scatter.

Model/MPD	**Valid** (%)	**Unique** (%)	**QED** > 0.6 (%)	**SA** < 5 (%)	**logP in range** (%)	**MPS** (%)	**GM**
Proposed/4	100	90.6	16.9 ± 0.5	61.3 ± 0.3	94.6 ± 0.6	11.9	1821
Proposed/4 (Top-1000)	100	86.3	100	92.4	99.7	92.1	921
InstGAN	97.7	98.7	93.0	99.0	89.0	–	–
L-MolGAN	97.5	–	86.1	55.0	16.2	–	–

Table 3: Comparison with baselines. Single-property satisfaction indicates the percentage of molecules meeting $\{\text{QED} > 0.6, \text{SA} < 5, \text{logP} \in [-0.5, 5]\}$ (Bickerton et al. 2012; Lipinski et al. 2012; Ertl and Schuffenhauer 2009). Multi-property satisfaction (MPS) reports the percentage and count of good molecules (GM) meeting all three constraints jointly. (-) means not reported.

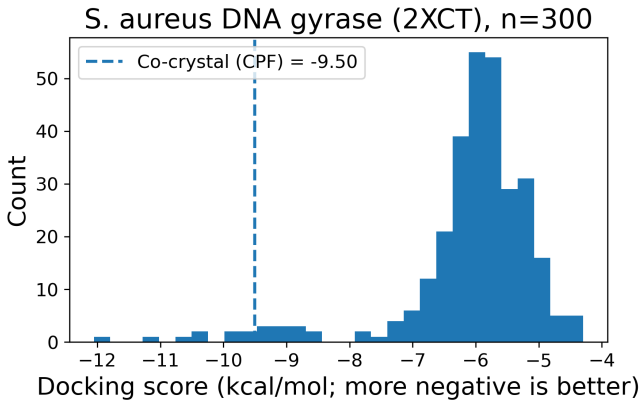

(a)

Co-crystal (ciprofloxacin)	-9.50
Mean score (n=300)	-6.10
# ligands $<$ co-crystal	9

(b)

Figure 8: (a) Docking score distribution. (b) Docking on *S. aureus*

timizing properties independently.

Conclusion

We introduce a descriptor-guided, patch-based GAN framework for antibiotic-like molecule generation. By aligning latent spaces with RDKit descriptors and enabling controllable synthesis through a modular pipeline, our method surfaces hundreds of structurally diverse molecules that jointly satisfy drug-likeness thresholds – a benchmark not addressed by prior models. Large-scale stress tests (50k samples) confirm resilience to mode collapse, while docking and QSAR analyses provide biological plausibility, yielding ligands that surpass co-crystal references and ciprofloxacin at validated antibacterial targets. This work demonstrates a scalable approach to interpretable, property-driven molecular design within clinically relevant domains. Future efforts will extend this framework to other therapeutic areas and refine reward formulations to better penalize synthetic complexity, bringing generated compounds closer to real antibiotics.

References

Andronov, M.; Andronova, N.; Wand, M.; Schmidhuber, J.; and Clevert, D.-A. 2025. Accelerating the inference of string generation-based chemical reaction models for industrial applications. *Journal of Cheminformatics*, 17(1): 31.

Bian, Y.; and Xie, X.-Q. 2022. Artificial intelligent deep learning molecular generative modeling of scaffold-focused and cannabinoid CB2 target-specific small-molecule sublibraries. *Cells*, 11(5): 915.

Bickerton, G. R.; Paolini, G. V.; Besnard, J.; Muresan, S.; and Hopkins, A. L. 2012. Quantifying the chemical beauty of drugs. *Nature chemistry*, 4(2): 90–98.

Cauchy, T.; Leguy, J.; and Da Mota, B. 2023. Definition and exploration of realistic chemical spaces using the connectivity and cyclic features of ChEMBL and ZINC. *Digital Discovery*, 2(3): 342–359.

Chen, L.; Yu, L.; and Gao, L. 2023. Potent antibiotic design via guided search from antibacterial activity evaluations. *Bioinformatics*, 39(2): btad059.

Daoud, N. E.-H.; Borah, P.; Deb, P. K.; Venugopala, K. N.; Hourani, W.; Alzweiri, M.; Bardaweel, S. K.; and Tiwari, V. 2021. ADMET Profiling in Drug Discovery and Development: Perspectives of In Silico, In Vitro and Integrated Approaches. *Current Drug Metabolism*, 22(7): 503–522.

De Cao, N.; and Kipf, T. 2018. MolGAN: An implicit generative model for small molecular graphs. In *ICML Workshop on Theoretical Foundations and Applications of Deep Generative Models*.

Du, Y.; Jamasb, A. R.; Guo, J.; Fu, T.; Harris, C.; Wang, Y.; Duan, C.; Liò, P.; Schwaller, P.; and Blundell, T. L. 2024. Machine learning-aided generative molecular design. *Nature Machine Intelligence*, 6(6): 589–604.

Ertl, P.; and Schuffenhauer, A. 2009. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *Journal of cheminformatics*, 1(1): 8.

Ferrara, F.; Castagna, T.; Pantolini, B.; et al. 2024. The challenge of antimicrobial resistance (AMR): current status and future prospects. *Naunyn-Schmiedeberg's Archives of Pharmacology*, 397: 9603–9615.

Garcia Satorras, V.; Hoogeboom, E.; Fuchs, F.; Posner, I.; and Welling, M. 2021. E (n) equivariant normalizing flows. *Advances in Neural Information Processing Systems*, 34: 4181–4192.

Gebauer, N.; Gastegger, M.; and Schütt, K. 2019. Symmetry-adapted generation of 3d point sets for the targeted discovery of molecules. *Advances in neural information processing systems*, 32.

Glavatskikh, M.; Leguy, J.; Hunault, G.; Cauchy, T.; and Da Mota, B. 2019. Dataset's chemical diversity limits the generalizability of machine learning predictions. *Journal of Cheminformatics*, 11(1): 69.

Guimaraes, G. L.; Sanchez-Lengeling, B.; Outeiral, C.; Farias, P. L.; and Aspuru-Guzik, A. 2017. Objective-Reinforced Generative Adversarial Networks (ORGAN) for Sequence Generation Models. *arXiv preprint arXiv:1705.10843*.

Haddad, R.; Litsa, E. E.; Liu, Z.; Yu, X.; Burkhardt, D.; and Bhisetti, G. 2025. Targeted molecular generation with latent reinforcement learning. *Scientific Reports*, 15(1): 15202.

Higgins, I.; Matthey, L.; Pal, A.; Burgess, C.; Glorot, X.; Botvinick, M.; Mohamed, S.; and Lerchner, A. 2017. beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. In *International Conference on Learning Representations (ICLR)*.

Hoogeboom, E.; Satorras, V. G.; Vignac, C.; and Welling, M. 2022. Equivariant diffusion for molecule generation in 3d. In *International conference on machine learning*, 8867–8887. PMLR.

Hu, C.; Li, S.; Yang, C.; Chen, J.; Xiong, Y.; Fan, G.; Liu, H.; and Hong, L. 2023. ScaffoldGVAE: scaffold generation and hopping of drug molecules via a variational autoencoder based on multi-view graph neural networks. *Journal of Cheminformatics*, 15(1): 91.

Jin, W.; Barzilay, R.; and Jaakkola, T. 2020. Hierarchical generation of molecular graphs using structural motifs. In *International conference on machine learning*, 4839–4848. PMLR.

Kingma, D. P.; Welling, M.; et al. 2013. Auto-encoding variational bayes.

Landrum, G. 2013. RDKit: Open-source cheminformatics. http://www.rdkit.org. Accessed: 2025-07-14.

Lipinski, C. A.; Lombardo, F.; Dominy, B. W.; and Feeney, P. J. 2012. Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. *Advanced drug delivery reviews*, 64: 4–17.

Luo, Y.; and Ji, S. 2022. An autoregressive flow model for 3d molecular geometry generation from scratch. In *International conference on learning representations (ICLR)*.

Maziarka, Ł.; Pocha, A.; Kaczmarczyk, J.; Rataj, K.; Danel, T.; and Warchoł, M. 2020. Mol-CycleGAN: a generative model for molecular optimization. *Journal of Cheminformatics*, 12(1): 2.

Preuer, K.; Renz, P.; Unterthiner, T.; Hochreiter, S.; and Klambauer, G. 2018. Fréchet ChemNet distance: a metric for generative models for molecules in drug discovery. *Journal of chemical information and modeling*, 58(9): 1736–1741.

Ramakrishnan, R.; Dral, P. O.; Rupp, M.; and von Lilienfeld, O. A. 2014. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific Data*, 1: 140022.

Rossen, L.; Sirockin, F.; Schneider, N.; and Grisoni, F. 2024. Scaffold Hopping with Generative Reinforcement Learning. *Journal of Chemical Information and Modeling*.

Schuh, M. G.; Hesse, J.; and Sieber, S. A. 2025. AI-guided antibiotic discovery pipeline from target selection to compound identification. *arXiv preprint arXiv:2504.11091*.

Shi, C.; Xu, M.; Zhu, Z.; Zhang, W.; Zhang, M.; and Tang, J. 2020. Graphaf: a flow-based autoregressive model for molecular graph generation. *arXiv preprint arXiv:2001.09382*.

Swanson, K.; Liu, G.; Catacutan, D. B.; Arnold, A.; Zou, J.; and Stokes, J. M. 2024. Generative AI for designing and validating easily synthesizable and structurally novel antibiotics. *Nature machine intelligence*, 6(3): 338–353.

Tang, H.; Li, C.; Kamei, S.; Yamanishi, Y.; and Morimoto, Y. 2024. Molecular Generative Adversarial Network with Multi-Property Optimization. *arXiv preprint arXiv:2404.00081*.

Tingle, B. I.; Tang, K. G.; Castanon, J. M.; et al. 2023. ZINC-22-A Free Multi-Billion-Scale Database of Tangible Compounds for Ligand Discovery. *Journal of Chemical Information and Modeling*, 63(4): 1166–1176.

Tong, X.; Liu, X.; Tan, X.; Li, X.; Jiang, J.; Xiong, Z.; Xu, T.; Jiang, H.; Qiao, N.; and Zheng, M. 2021. Generative models for de novo drug design. *Journal of Medicinal Chemistry*, 64(19): 14011–14027.

Tsujimoto, Y.; Hiwa, S.; Nakamura, Y.; Oe, Y.; and Hiroyasu, T. 2021. L-MolGAN: An improved implicit generative model for large molecular graphs. *Chemrxiv*.

Voloboev, S. 2024. A Review on Fragment-based De Novo 2D Molecule Generation. *arXiv preprint arXiv:2405.05293*.

Wei, L.; Fu, N.; Song, Y.; Wang, Q.; and Hu, J. 2023. Probabilistic generative transformer language models for generative design of molecules. *Journal of Cheminformatics*, 15(88).

WHO. 2022. Lack of innovation set to undermine antibiotic performance and health gains. https://www.who.int/news/item/22-06-2022-22-06-2022-lack-of-innovation-set-to-undermine-anti\\biotic-performance-and-health-gains.

Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How Powerful are Graph Neural Networks? In *International Conference on Learning Representations (ICLR)*.

You, J.; Liu, B.; Ying, Z.; Pande, V.; and Leskovec, J. 2018. Graph convolutional policy network for goal-directed molecular graph generation. *Advances in neural information processing systems*, 31.

Zdrazil, B.; Felix, E.; Hunter, F.; Manners, E. J.; Blackshaw, J.; Corbett, S.; de Veij, M.; Ioannidis, H.; Mendez Lopez, D.; Mosquera, J. F.; et al. 2023. The ChEMBL Database in 2023: a drug discovery platform spanning multiple bioactivity data types and time periods. *Nucleic Acids Research*, 52(D1): D930–D940.