

The Eurasia Proceedings of Science, Technology, Engineering and Mathematics (EPSTEM), 2025

Volume 38, Pages 512-524

IConTES 2025: International Conference on Technology, Engineering and Science

AI-Driven Molecular Design: Synergizing Deep Generative Models with Evolutionary Optimization

Abbad Houda

Djillali Liabes University of Sidi Bel Abbes

Taieb Brahim Mohammed

Djillali Liabes University of Sidi Bel Abbes

Oulladji Latefa

Djillali Liabes University of Sidi Bel Abbes

Abbad Leila

ENSTA High School

Abstract: Artificial intelligence (AI) is reshaping drug discovery by enabling efficient and precise identification of novel therapeutics. This review examines the synergistic use of deep generative models, such as Long Short-Term Memory (LSTM) networks, Variational Autoencoders (VAEs), and Graph Attention Networks (GATs), together with evolutionary optimization techniques, including genetic algorithms and multi-objective evolutionary strategies. While deep learning architectures excel at capturing complex molecular representations and generating chemically valid compounds, evolutionary algorithms provide complementary strengths in global exploration and multi-objective trade-off optimization. The combination of these two paradigms offers a powerful and complementary toolkit: deep learning provides the capacity to learn rich chemical features and propose innovative scaffolds, whereas evolutionary methods ensure efficient navigation of chemical space and balanced optimization across multiple drug-like criteria. Through comparative analyses, quantitative benchmarks, and illustrative figures, we highlight how integrating generative and evolutionary paradigms can accelerate de novo molecular design, reduce development timelines, and lower costs. We also address technical and ethical challenges. In particular, our ongoing research explores hybrid frameworks that combine variational autoencoders, graph neural predictors, Colibri algorithm and Genetic algorithms with fragment-based crossover, and dynamic multi-objective penalties to further enhance chemical validity, pharmacological relevance, and synthetic accessibility. Future efforts aim to demonstrate that such hybrid frameworks can bridge the gap between theoretical innovation and practical drug development, bringing AI-driven discovery closer to real-world therapeutic breakthroughs.

Keywords: Drug discovery, Molecule design, Deep generative models, Evolutionary algorithms, Multi-objective optimization.

Introduction

Drug discovery is a cornerstone of biomedical innovation but remains an extremely costly and time-consuming endeavor, often requiring over a decade and billions of dollars to bring a single drug to market (DiMasi, Grabowski, & Hansen, 2016). Despite advances in computational chemistry and high-throughput screening, attrition rates in clinical trials remain high due to issues of efficacy, toxicity, and poor pharmacokinetics (Paul et al., 2010). These challenges underscore the urgent need for more efficient strategies capable of exploring the

- This is an Open Access article distributed under the terms of the Creative Commons Attribution-Noncommercial 4.0 Unported License, permitting all non-commercial use, distribution, and reproduction in any medium, provided the original work is properly cited.

- Selection and peer-review under responsibility of the Organizing Committee of the Conference

vast chemical space, estimated to contain up to 10^{60} drug-like molecules, while simultaneously optimizing multiple therapeutic objectives (Polishchuk, Madzhidov, & Varnek, 2013).

Conventional in silico approaches, such as docking and quantitative structure–activity relationship (QSAR) models, provide valuable predictive tools but often suffer from limited generalizability, high dependence on prior knowledge, and poor scalability when faced with unprecedented chemical scaffolds (Schneider et al., 2020). Similarly, in vitro assays, although indispensable for experimental validation, are labor-intensive, expensive, and inherently sequential, restricting their ability to explore chemical diversity at scale (Vamathevan et al., 2019). The fragmentation between computational predictions and experimental testing continues to delay the discovery of effective therapeutics, highlighting the need for integrated frameworks that are both exploratory and optimization-oriented.

Recent progress in artificial intelligence (AI) has introduced generative models, such as variational autoencoders (VAEs), generative adversarial networks (GANs), and graph neural networks (GNNs), which can learn complex molecular distributions and design novel compounds with promising drug-like properties (Zhavoronkov et al., 2019; Blaschke et al., 2020). These models excel in capturing latent chemical features and generating structurally valid molecules, but often struggle with balancing multiple objectives, such as potency, selectivity, and synthetic accessibility. Evolutionary algorithms (EAs), by contrast, are particularly effective at global exploration and multi-objective optimization, using mechanisms like mutation, crossover, and selection to iteratively refine candidate solutions (Brown et al., 2019). Integrating deep generative models with EAs offers a synergistic paradigm: the former drives innovation by proposing diverse molecular scaffolds, while the latter ensures efficient navigation of chemical space and optimization across competing therapeutic criteria (Nigam et al., 2022). This hybridization promises not only to accelerate de novo molecular design but also to reduce attrition rates and shorten development timelines, ultimately transforming AI-driven drug discovery.

The remainder of this paper is organized as follows. Section 2 provides an overview of deep generative models for molecular design, including recurrent neural networks, variational auto-encoders, generative adversarial networks, and graph-based models, highlighting their respective strengths and limitations. Section 3 discusses evolutionary optimization methods, with an emphasis on genetic algorithms, multi-objective strategies, and chemically valid genetic operators tailored to drug discovery. Section 4 explores the synergy between generative models and evolutionary algorithms, presenting hybrid frameworks, case studies, and benchmark comparisons, as well as illustrative workflows. Section 5 outlines future research directions, focusing on dynamic multi-objective penalties, integration with reinforcement learning, and the development of fully automated, closed-loop drug discovery pipelines. In addition, Section 6 highlights our contributions to AI-Driven drug discovery with novel methodological innovations. Finally, Section 7 concludes the paper with a synthesis of key insights and perspectives on the role of hybrid generative–evolutionary paradigms in accelerating AI-driven molecular design.

Deep Generative Models for Molecular Design

The emergence of deep generative models has profoundly transformed de novo molecular design, enabling the automated generation of structurally valid and pharmacologically relevant compounds (Table. 1). By learning latent representations of chemical space, these models provide scalable frameworks to design molecules with desired physicochemical, biological, and pharmacokinetic properties. The following subsections review the most prominent generative architectures employed in molecular design.

Recurrent Neural Networks (RNNs, LSTM, GRU)

Recurrent neural networks (RNNs) and their gated variants: Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRU), have been widely adopted for sequence-based molecular generation, typically using SMILES (Simplified Molecular Input Line Entry System) representations (Olivecrona et al., 2017). These models leverage sequential dependencies to learn syntax and semantics of chemical strings, ensuring the generation of syntactically valid SMILES. LSTMs are particularly effective at capturing long-range dependencies, reducing the risk of invalid outputs, while GRUs provide computational efficiency without significantly compromising accuracy (Gupta et al., 2018). However, SMILES-based models remain sensitive to syntax errors, and their reliance on a single linear representation of molecules can limit structural diversity.

Variational Autoencoders (VAEs)

Variational autoencoders (VAEs) introduced probabilistic latent representations into molecular design, enabling interpolation and continuous optimization in chemical space (Kingma & Welling, 2014). In this framework, molecules are encoded into a latent vector space and decoded back into valid molecular structures, which facilitates property-driven optimization via gradient-based methods (Gómez-Bombarelli et al., 2018). VAEs can also incorporate chemical constraints and multi-objective loss functions to balance validity, novelty, and synthetic accessibility (Jin et al., 2018). Nevertheless, VAEs often suffer from posterior collapse and low reconstruction fidelity, particularly for complex molecular graphs, which has motivated hybridization with graph neural networks and reinforcement learning strategies (Winter et al., 2019).

Generative Adversarial Networks (GANs)

Generative adversarial networks (GANs) employ a competitive setup between a generator, which proposes candidate molecules, and a discriminator, which distinguishes between real and generated samples (Goodfellow et al., 2014). In molecular design, GANs have demonstrated strong potential for producing chemically diverse scaffolds while aligning outputs with drug-like distributions (Kadurin et al., 2017). Advanced implementations, such as ORGAN (Objective-Reinforced GAN), integrate reinforcement learning to steer generation toward molecules with optimized bioactivity (Guimaraes et al., 2017). Despite their promise, GANs remain difficult to train due to instability, mode collapse, and the challenge of ensuring strict chemical validity.

Graph Neural Networks (GNNs, GATs)

Graph neural networks (GNNs) represent molecules as graphs, where atoms correspond to nodes and bonds to edges, allowing direct learning from molecular topology (Gilmer et al., 2017). Variants such as Graph Convolutional Networks (GCNs) and Graph Attention Networks (GATs) enhance molecular representation by capturing both local and global structural dependencies (Velickovic et al., 2018). Recent advances in graph-based generative models, such as GraphVAE and GraphAF, enable the direct generation of molecular graphs, bypassing SMILES limitations (You et al., 2018; Shi et al., 2020). These models demonstrate superior performance in terms of validity, novelty, and scaffold diversity. However, challenges remain in balancing efficiency with expressiveness, as graph-based generation is computationally intensive and requires sophisticated decoding strategies.

Figure 1 summarizes the major classes of deep generative models applied in de novo molecular design. Each architecture provides a distinct strategy for navigating chemical space and balancing structural validity, novelty, and property optimization.




Figure 1. Overview of deep generative architectures for molecular design

Strengths and Limitations of Generative Models

Generative models collectively provide a paradigm shift in drug discovery by automating scaffold innovation, accelerating chemical exploration, and supporting multi-property optimization (see Table 1). RNNs are simple and efficient but constrained by SMILES syntax. VAEs enable smooth latent-space optimization but struggle with reconstruction fidelity. GANs encourage diversity and novelty but suffer from instability and mode collapse. GNNs and GATs provide chemically faithful graph-based generation but demand high computational resources and complex training strategies. In practice, these strengths and limitations suggest that no single generative model is universally optimal. Instead, hybrid frameworks, combining sequence-based, probabilistic, adversarial, and graph-based strategies, are increasingly recognized as essential for robust de novo molecular design (Brown et al., 2019; Zhavoronkov et al., 2019).

Table 1. Comparative overview of generative models

Generative Model	RNNs (LSTM/GR)	VAEs	GANs	GNN-based Models
Key References	Gupta et al. (2018); Olivecrona et al. (2017)	Kingma & Welling (2014); Gómez-Bombarelli et al. (2018)	Goodfellow et al. (2014); Guimaraes et al. (2017)	Gilmer et al. (2017); Jin et al. (2018)
Representation	Sequential SMILES	Latent space (SMILES/graphs)	SMILES or latent embeddings	Graphs (atoms & bonds)
Strengths	Simple; RL-friendly	Smooth latent space; interpolation	Models complex distributions; reinforcement integration	Captures chemical validity; strong for property prediction
Limitations	Sensitive to SMILES syntax; limited space coverage	Invalid molecules; limited diversity	Training instability; mode collapse	Computationally expensive
Implementations	REINVENT; DeepChem modules	ChemicalVAE; MolecularRNN-VAE	ORGAN; MoleculeGAN; druGAN	JT-VAE; GraphVAE; MolGAN; GCPN

Evolutionary Optimization in Drug Discovery

Evolutionary optimization encompasses a family of population-based metaheuristic methods inspired by natural evolution (mutation, recombination, selection) (Deb et al., 2002a). In molecular design, these methods are used both as standalone search engines and as complementing modules to deep generative models, enabling explicit optimization of multiple competing objectives (e.g., potency, selectivity, ADMET, synthetic accessibility) while preserving structural diversity and interpretability (Schneider, 2018; Zhong et al., 2019). Recent advances have integrated graph-based representations, fragment-level operators, and efficient multi-objective schemes to make evolutionary approaches increasingly practical for de novo drug design (Polishchuk et al., 2013; Nigam et al., 2022; Deb et al., 2002b).

Principles of Genetic Algorithms

Genetic algorithms (GAs) are among the most widely applied evolutionary methods in computational chemistry. A typical GA for molecular design maintains a population of candidate molecules that are iteratively evolved using operators that mimic biological evolution: selection (based on a fitness function), crossover (recombination of parental substructures), and mutation (random local changes) (Mitchell, 1998; Holland, 1992). The fitness function can be single-objective (e.g., predicted binding affinity) or compositional (weighted combination of multiple properties). Key algorithmic choices include representation (SMILES, graph, fragment-based encoding), the encoding of genotype-to-phenotype mappings, population sizing, and selection pressure (Schneider, 2018). Representations that embed chemical knowledge (e.g., fragment- or graph-based encodings) reduce the generation of chemically invalid offspring and speed convergence (Polishchuk et al., 2013). Practical GA pipelines often use surrogate predictors (ML models) to evaluate fitness cheaply, interleaving costly physics-based or experimental evaluations only for top candidates (Jin et al., 2020; Zhou et al., 2019).

Genetic algorithms encode molecules as chromosomes, commonly represented by SMILES strings, graphs, or molecular fingerprints, and evolve them toward improved fitness. Each candidate solution $x_i \in X$ is evaluated through a fitness function:

$$F(x_i) = \sum_{i=1..K} (w_k \times f_k(x_i))$$

where $f_k(x_i)$ represents the k -th molecular property (e.g., binding affinity, solubility, toxicity), and w_k denotes the weight assigned to this objective (Nigam et al., 2021). Selection, crossover, and mutation operators are applied iteratively until convergence. This makes GAs particularly suitable for navigating discontinuous, high-dimensional, and multi-modal chemical spaces (Brown et al., 2019).

To better illustrate the operational cycle of genetic algorithms in molecular design, Figure 2 presents a schematic workflow highlighting the main evolutionary steps. Starting from an initial population of candidate molecules encoded as SMILES strings or graph-based representations, the algorithm iteratively applies selection, crossover, and mutation operators to generate new offspring. A fitness evaluation module, often supported by surrogate machine learning models or physics-based simulations, guides the evolutionary search toward improved molecular properties such as binding affinity, solubility, or synthetic accessibility (Nigam et al., 2021; Jensen, 2019). The iterative nature of this workflow enables the progressive refinement of molecular candidates while preserving chemical diversity and avoiding premature convergence (Zhavoronkov et al., 2019).




Figure 2. Workflow of a genetic algorithm in de novo molecular design

Multi-objective Evolutionary Strategies

Drug discovery is inherently a multi-objective optimization problem, where candidate molecules must balance several often conflicting criteria, such as potency, selectivity, solubility, toxicity, and synthetic accessibility. Unlike single-objective optimization methods, multi-objective evolutionary algorithms (MOEAs) are particularly well-suited to this task because they do not aim to find a single best solution but rather approximate the Pareto front, a set of trade-off solutions where improving one objective would deteriorate at least one other (Deb et al., 2002; Zhang & Li, 2007).

Formally, the optimization problem can be expressed as:

$$\text{Optimize } F(x) = (f_1(x), f_2(x), \dots, f_k(x))$$

subject to chemical validity constraints (e.g., valence rules, synthetic feasibility). Here, $f_k(x)$ represents the k^{th} molecular property of candidate solution x , and the goal is to evolve a population of molecules such that their fitness values collectively approximate the Pareto-optimal front (Nigam et al., 2021).

Among MOEAs, NSGA-II (Non-dominated Sorting Genetic Algorithm II) is the most widely applied in cheminformatics due to its simplicity, efficiency, and ability to preserve population diversity. It relies on non-dominated sorting to classify solutions into Pareto fronts and uses a crowding distance metric to ensure spread across the front. This makes NSGA-II particularly effective for drug design tasks such as optimizing potency against synthetic accessibility or ADMET properties (Stahl et al., 2022).

Table 2. Comparative overview of NSGA-II and MOEA/D in Multi-Objective Molecular Design

Criterion	NSGA-II	MOEA/D	References
Optimization paradigm	Non-dominated sorting with crowding-distance to rank and preserve Pareto diversity	Decomposition of multi-objective problem into scalar subproblems optimized cooperatively	Deb et al., 2002; Zhang & Li, 2007
Diversity preservation	Ensures diversity using crowding-distance and elitism	Maintains diversity through neighborhood relations in decomposed subproblems	Deb et al., 2002; Li & Zhang, 2009
Scalability	Effective for 2–3 objectives but performance degrades with higher dimensions	More scalable for many-objective problems due to decomposition strategy	Ishibuchi et al., 2016
Convergence	Strong convergence for low-dimensional objectives, sensitive to parameter tuning	Competitive convergence in high-dimensional settings; less sensitive to number of objectives	Zhang & Li, 2007; Ishibuchi et al., 2016
Computational complexity	$O(MN^2)$ per generation (M = objectives, N = population size)	$O(NT)$ per generation (T = neighborhood size), often lower computational overhead	Deb et al., 2002; Zhang & Li, 2007
Applications in drug discovery	Used for optimizing potency vs. synthetic accessibility; effective in generating diverse candidate molecules	Applied to multi-objective molecular generation (e.g., ADMET balance); advantageous for complex property landscapes	Nigam et al., 2021; Gao et al., 2022
Limitations	May lose performance in high-dimensional spaces; crowding-distance can bias diversity	Requires careful decomposition strategy; neighborhood definition affects results	Deb et al., 2002; Zhang & Li, 2007; Ishibuchi et al., 2016

Alternatively, MOEA/D (Multi-Objective Evolutionary Algorithm based on Decomposition) decomposes a multi-objective problem into a set of scalar optimization subproblems, each optimized in parallel while sharing information with neighboring subproblems. This decomposition approach often provides improved scalability and convergence for high-dimensional objective spaces (Zhang & Li, 2007). Recent studies in molecular optimization suggest that MOEA/D variants can perform comparably, or even better, than NSGA-II in specific

bi-objective and tri-objective settings (Brown et al., 2019). Finally, recent adaptations of MOEAs in drug discovery incorporate dynamic penalty strategies to account for drug-likeness and synthetic accessibility, as well as graph-based molecular encodings to ensure the chemical validity of offspring during evolution (Zhavoronkov et al., 2019; Stahl et al., 2022). These improvements make evolutionary multi-objective optimization a powerful and increasingly practical tool for guiding de novo molecular design.

Although both NSGA-II and MOEA/D have demonstrated strong performance in molecular multi-objective optimization, they rely on fundamentally different paradigms, non-dominated sorting versus decomposition. Table 2 summarizes their comparative characteristics, offering insights into when each algorithm may be preferable in the context of drug discovery.

Fragment-Based and Crossover Operators in Chemical Space

To ensure the generation of chemically valid offspring during evolutionary optimization, crossover and mutation operators are frequently constrained by fragment-based recombination strategies (Brown et al., 2019). Unlike naive string-level manipulations of SMILES that can easily produce invalid structures, fragment-based approaches operate on chemically meaningful units, such as rings, linkers, or functional groups. For two parent molecules x_p and x_q , a fragment-based crossover operator C can be formally defined as:

$$x_{child} = C(x_p, x_q) = \text{merge}(\text{frag}(x_p), \text{frag}(x_q))$$

where $\text{frag}(\cdot)$ denotes the extraction of synthetically feasible substructures. This ensures that the recombination process preserves valence rules, avoids bond-breaking artifacts, and yields molecules that are more likely to be chemically stable and synthetically accessible (Polishchuk et al., 2013).

Mutation operators within this framework typically involve localized structural perturbations such as functional group substitutions, ring contractions or expansions, and stereochemical inversions. These operations allow the exploration of diverse regions of chemical space while maintaining molecular validity (Nigam et al., 2020). Importantly, chemically aware operators significantly reduce the occurrence of invalid SMILES strings, a major drawback in purely syntax-based generative approaches, and contribute to the design of molecules that are more readily synthesizable (Yoshimori et al., 2021; Segler et al., 2018).

In this way, fragment-based crossover and chemically guided mutation form a core component of evolutionary molecular design pipelines, balancing the trade-off between exploration of novel chemical scaffolds and exploitation of synthetically tractable structures.

GA Implementations in Drug Discovery

The application of genetic algorithms (GAs) in drug discovery has evolved significantly, moving from early SMILES-based encodings to sophisticated graph and latent space representations. Several notable implementations illustrate this progression, each addressing specific challenges in de novo molecular design. MolEvol represents one of the earliest GA-based molecular design frameworks, relying on SMILES string encodings to explore chemical space (Brown et al., 2019). Its optimization objectives included binding affinity and drug-likeness quantified by the Quantitative Estimate of Drug-likeness (QED), while operators such as fragment-based crossover and mutations were employed to preserve chemical validity. Despite its simplicity, MolEvol demonstrated the feasibility of evolutionary search in chemical space.

In contrast, JANUS introduced a graph-based encoding scheme that enabled a more chemically intuitive representation of molecules (Nigam et al., 2021). By incorporating adaptive crossover and mutation strategies, JANUS efficiently balanced multiple objectives such as potency, solubility, and toxicity. This multi-objective optimization framework addressed the inherent trade-offs in drug design, showcasing the power of GAs in multi-criteria decision-making.

MolGA further refined graph-based representations by implementing graph crossover and edge mutation operators tailored to molecular structures (Zhou et al., 2019). This design ensured that generated molecules adhered to chemical validity constraints while maintaining high diversity. MolGA's emphasis on synthesizability and drug-likeness marked an important step toward practical drug discovery applications.

The integration of genetic algorithms with deep generative models is exemplified by DEFactor, which leveraged latent space embeddings as an encoding scheme (Assouel et al., 2018). In this approach, crossover operations were performed in the learned latent space, enabling smooth exploration of molecular manifolds. This hybridization of GAs with deep learning allowed optimization based on docking scores and QED, bridging symbolic evolutionary search with representation learning.

Finally, Hybrid GA-ML approaches combined graph embeddings with machine learning predictors to optimize pharmacokinetic and toxicological endpoints (Stahl et al., 2022). By reinforcing crossover and guiding mutations with predictive models, these methods reduced the computational cost associated with expensive molecular evaluations such as docking or quantum chemistry. Such hybrid frameworks exemplify the shift toward efficiency-driven GA implementations in modern drug discovery pipelines.

Overall, these implementations highlight the adaptability of genetic algorithms to various molecular representations and optimization objectives. From string-based encodings to latent space manipulations, GAs continue to provide a versatile search paradigm for balancing the competing objectives of potency, safety, and synthesizability in drug discovery.

Limitations and Challenges

Although evolutionary methods bring transparency and multi-objective capability, they face several challenges in practice:

- **Computational Cost:** Populations and many generations require large numbers of fitness evaluations; expensive physics-based or experimental scoring limits throughput. Surrogate models and active sampling have been proposed to mitigate this cost, but they introduce model-bias risks (Nigam et al., 2020; Stahl et al., 2022).
- **Validity vs. Novelty Trade-off:** Aggressive mutation and crossover increase novelty but often yield chemically invalid or synthetically infeasible molecules. Balancing exploration and exploitation is therefore nontrivial. Fragment-aware encodings and chemically valid operators help reduce invalid candidates but do not eliminate the problem entirely (Brown et al., 2019; Yoshimori et al., 2021).
- **Objective Modeling Errors:** Machine learning surrogates can misrank candidates, and noisy predictors may drive premature convergence toward artifacts. Robust uncertainty quantification and multi-fidelity evaluation pipelines are recommended (Vamathevan et al., 2019; Schneider et al., 2020).
- **Scalability to Many Objectives:** Many-objective optimization (>4–5 objectives) complicates Pareto selection and diversity maintenance. Specialized algorithms such as NSGA-II or decomposition-based methods like MOEA/D have been adapted to molecular discovery to address this issue (Deb et al., 2002; Zhang & Li, 2007; Ishibuchi et al., 2016).
- **Integration with Wet-Lab Workflows:** Translating evolutionary outputs into synthesizable, testable molecules requires retrosynthetic planning, procurement considerations, and assay translation areas where further automation and standardization are needed (Segler et al., 2018; Zhavoronkov et al., 2019).

Overall, evolutionary optimization remains a powerful and interpretable approach for multi-objective molecular design, particularly when combined with modern graph representations, fragment-aware operators, and surrogate models. The trend in recent literature is toward hybrid pipelines that combine the generative power of deep models with the explicit search and control offered by evolutionary algorithms, yielding pragmatic systems capable of producing chemically plausible, multi-objective-optimized candidates at scale (Gómez-Bombarelli et al., 2018; Nigam et al., 2021).

Synergizing Generative Models with Evolutionary Algorithms

Complementary Strengths of the Two Paradigms

Generative models and evolutionary algorithms (EAs) offer distinct yet complementary advantages for molecular design (refer to Table 3). Generative models, such as variational autoencoders (Kingma & Welling, 2014; Gómez-Bombarelli et al., 2018), generative adversarial networks (Goodfellow et al., 2014; Kadurin et al., 2017; Guimaraes et al., 2017), and reinforcement learning-based graph models (Olivecrona et al., 2017; You et al., 2018; Shi et al., 2020), excel at capturing the complex distribution of chemical space and generating novel

structures. However, these models may suffer from biases introduced by training data and may not directly optimize multiple drug-relevant objectives (Schneider, 2018; Zhong et al., 2019).

In contrast, evolutionary algorithms, particularly multi-objective methods such as NSGA-II (Deb et al., 2002) and MOEA/D (Zhang & Li, 2007), are inherently suited for balancing competing objectives like potency, solubility, and toxicity (Nigam et al., 2021). EAs provide explicit control over diversity maintenance and can incorporate chemically valid crossover and mutation operators (Yoshimori et al., 2021). By combining these paradigms, generative models can propose diverse and syntactically valid candidates, while EAs refine and optimize them in accordance with medicinal chemistry constraints. Table 3 highlights the complementarity of generative models and genetic algorithms:

Table 3: Comparative analysis of deep generative models vs. evolutionary algorithms in molecular design

Criterion	Deep Generative Models	Evolutionary Algorithms	Benchmark References
Validity	70–99% >99.9% (SELFIES:)	100% (fragment constraints)	Krenn et al. (2020); Segler et al. (2018)
Synthesizability (S A score ≤4)	15–40%	75–92%	Gao & Coley (2020); Nigam et al. (2019)
Novelty (Tanimoto <0.4)	80–95%	45–70%	Brown et al. (2019); You et al. (2018)
Multi-Objective Support	Limited by reward design	Native integration via fitness	Li et al. (2018); Steinmann & Jensen (2021)

Hybrid framework

Hybrid frameworks aim to leverage the expressive generation capabilities of deep models with the robust search and optimization of evolutionary algorithms. One common strategy is to embed molecules in continuous latent spaces learned by autoencoders or graph neural networks (Blaschke et al., 2020; Winter et al., 2019), where genetic operators such as crossover and mutation can be applied more smoothly (Assouel et al., 2018). Another approach integrates reinforcement learning-based generative models with EA-driven population management, ensuring balance between exploration and exploitation (Nigam et al., 2020; Stahl et al., 2022).

Such frameworks allow iterative interplay: generative models explore wide areas of chemical space, while evolutionary operators guide optimization toward Pareto-optimal sets with respect to pharmacokinetics, toxicity, and synthetic accessibility (Ishibuchi et al., 2016; Schneider et al., 2020). This synergy directly addresses the limitations of purely generative approaches, particularly in multi-objective settings where trade-offs are unavoidable.

Case Studies and Benchmarks

Several implementations exemplify the synergy between generative and evolutionary paradigms (refer to Figure 3). The JANUS framework employs a parallel-tempered genetic algorithm guided by neural networks, enabling efficient multi-objective optimization of drug-like molecules (Nigam et al., 2021; Nigam et al., 2022). STONED, based on the SELFIES molecular representation, achieves rapid traversal and exploration of chemical space, demonstrating the effectiveness of EA-guided search without the need for large-scale training (Nigam, Pollice et al., 2022).

Benchmarking platforms such as GuacaMol (Brown et al., 2019) and comparative studies (Polishchuk et al., 2013) provide standardized metrics for evaluating generative-EA hybrids across objectives like quantitative estimate of drug-likeness (QED), binding affinity, and synthesizability. These studies confirm that hybrid methods often outperform standalone generative models or purely evolutionary searches in terms of chemical validity, novelty, and multi-objective performance. Figure 2 illustrates a typical workflow for hybrid generative–evolutionary drug design begins with a generative model trained on large chemical datasets (Gómez-Bombarelli et al., 2018; Gilmer et al., 2017). Candidate molecules generated in SMILES or graph form are

passed to an evolutionary optimization layer, where crossover, mutation, and selection are applied under multi-objective constraints (Deb et al., 2002; Li & Zhang, 2009). Feedback from predictive models (Vamathevan et al., 2019; Zhavoronkov et al., 2019) or docking simulations (Zhou et al., 2019) is used to update both the evolutionary fitness functions and the generative model parameters. Figure 3 illustrates the synergy between generative models and evolutionary systems for multi-objective molecular design:




Figure 3. Workflow of hybrid generative–evolutionary frameworks for multi-objective molecular design

Future Directions

The convergence of generative modeling and evolutionary optimization is still at an early stage, and several avenues remain open for advancing the field.

Dynamic Multi-Objective Penalties

Traditional multi-objective optimization often treats all objectives with static weights or Pareto dominance. However, drug discovery objectives such as potency, selectivity, and pharmacokinetic properties, rarely remain fixed throughout a project's lifecycle. Dynamic penalty schemes, in which weights adapt based on progress or project stage, could enable more pragmatic optimization. Such adaptive formulations would allow algorithms to progressively shift emphasis from chemical diversity toward synthesis feasibility and clinical relevance (Deb et al., 2002; Li & Zhang, 2009).

Integration with Reinforcement Learning

While evolutionary algorithms excel in population-based exploration, reinforcement learning (RL) methods provide fine-grained control over sequential molecular construction (Olivecrona et al., 2017; Zhou et al., 2019). Combining these paradigms could yield synergistic frameworks where RL policies guide local exploration, while evolutionary operators maintain population-level diversity and multi-objective balance. Such hybridization is especially promising for tasks requiring long-horizon credit assignment, such as scaffold hopping or synthesizability-aware optimization.

Towards Fully Automated, Closed-Loop Drug Discovery Pipelines

The ultimate goal of this research trajectory is to achieve autonomous design–make–test–analyze (DMTA) cycles (Schneider, 2018; Segler et al., 2018). Closing the loop between in silico design, robotic synthesis, and high-throughput screening will require robust interfacing of generative–evolutionary models with retrosynthetic planning engines (Segler et al., 2018) and laboratory automation platforms. Active learning and uncertainty-aware surrogate models (Vamathevan et al., 2019) will be critical for prioritizing experiments and reducing costs in these pipelines. Progress in this direction could enable self-driving laboratories that iteratively refine molecular candidates with minimal human intervention, accelerating the path from hypothesis to validated lead.

Our Contributions to AI-Driven Drug Discovery

Over the three past years, our research group has made several contributions to the rapidly evolving field of AI-driven drug discovery. Building on the synergy between deep generative models and evolutionary optimization, we have focused on developing deep generative and hybrid frameworks that combine data-driven molecular generation with multi-objective search strategies. These efforts aim to accelerate de novo molecular design, improve chemical validity, and enhance the interpretability of AI-assisted drug development.

Our first contribution Oulladji et al., 2025, explores the use of artificial intelligence and machine learning techniques in the drug discovery process, focusing on how data-driven models can accelerate target identification, molecular design, and compound optimization. The study highlights the importance of predictive modeling, deep learning, and virtual screening in reducing cost and time in drug development. Finally, it emphasizes future directions toward automated and intelligent drug design pipelines. Second research, Taieb Brahim et al., currently under peer review, proposes a hybrid framework that integrates a Variational Autoencoder (VAE) to generate diverse anticancer molecular structures and a multi-objective Genetic Algorithm to optimize them. The optimization process simultaneously targets predicted anticancer activity (GCN), drug-likeness (QED), synthetic accessibility (SA), and Lipinski compliance. Overall, the proposed approach efficiently explores chemical space and identifies promising, realistic, and synthesizable drug candidates.

Finally, our current contribution, Abbad et al., investigates the integration of deep generative models with a dual optimization strategy leveraging Genetic Algorithms (GA) and the Colibri algorithm. The generative model provides diverse initial molecular candidates, while the GA–Colibri optimization jointly refines them with respect to predicted bioactivity, drug-likeness, synthetic accessibility, and pharmacokinetic constraints. The designed approach accelerates de novo molecule discovery while maintaining a balance between potency, feasibility, and developability. Together, these contributions illustrate our commitment to advancing the intersection of artificial intelligence, cheminformatics, and computational drug design, paving the way toward more efficient, interpretable, and autonomous discovery workflows.

Conclusion

The integration of generative models with evolutionary algorithms is reshaping molecular design by uniting data-driven creativity with interpretable, multi-objective optimization. Emerging hybrid frameworks not only enhance chemical validity and synthesizability but also pave the way toward autonomous, closed-loop discovery pipelines. As adaptive optimization strategies and reinforcement learning integration mature, these approaches hold the potential to accelerate drug discovery dramatically, bridging the gap between theoretical innovation and tangible therapeutic breakthroughs. Ultimately, hybrid generative–evolutionary frameworks may enable the transition from exploratory molecular design to truly autonomous drug discovery.

Scientific Ethics Declaration

* The authors declare that the scientific ethical and legal responsibility of this article published in EPSTEM journal belongs to the authors.

Conflict of Interest

* The authors declare that they have no conflicts of interest

Funding

* This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Acknowledgements or Notes

* This article was presented as an oral presentation at the International Conference on Technology, Engineering and Science (www.icontes.net) held in Antalya/Türkiye on November 12-15, 2025.

References

Assouel, R., Ahmed, N., Wolf, G., Bresson, X., & Cohen, J. (2018). DEFactor: Differentiable edge factorization-based probabilistic graph generation. *arXiv*.

Bagal, V., Aggarwal, R., Vinod, P. K., & Priyakumar, U. D. (2021). MolGPT: Molecular generation using a transformer-decoder model. *Journal of Chemical Information and Modeling*, 61(12), 5804–5814.

Blaschke, T., Olivecrona, M., Engkvist, O., Bajorath, J., & Chen, H. (2020). Application of generative autoencoder in de novo molecular design. *Molecular Informatics*, 39(1-2), 1900123.

Brown, N., Fiscato, M., Segler, M. H. S., & Vaucher, A. C. (2019). GuacaMol: Benchmarking models for de novo molecular design. *Journal of Chemical Information and Modeling*, 59(3), 1096-1108.

Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2), 182-197.

Deb, K., Thiele, L., Laumanns, M., & Zitzler, E. (2002). Scalable multi-objective optimization test problems. *IEEE Congress on Evolutionary Computation*, 825-830.

DiMasi, J. A., Grabowski, H. G., & Hansen, R. W. (2016). Innovation in the pharmaceutical industry: New estimates of R&D costs. *Journal of Health Economics*, 47, 20-33.

Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., & Dahl, G. E. (2017). Neural message passing for quantum chemistry. *ICML*, 1263-1272.

Gómez-Bombarelli, R., Wei, J. N., Duvenaud, D., Hernández-Lobato, J. M., Sánchez-Lengeling, B., Sheberla, D., & Aspuru-Guzik, A. (2018). Automatic chemical design using a data-driven continuous representation of molecules. *ACS Central Science*, 4(2), 268-277.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., & Bengio, Y. (2014). Generative adversarial nets. *NeurIPS*, 2672-2680.

Guimaraes, G. L., Sanchez-Lengeling, B., Outeiral, C., Farias, P. L. C., & Aspuru-Guzik, A. (2017). ORGAN: Objective-reinforced generative adversarial networks. *arXiv*.

Gupta, A., Müller, A. T., Huisman, B. J., Fuchs, J. A., Schneider, P., & Schneider, G. (2018). Generative recurrent networks for de novo drug design. *Molecular Informatics*, 37(1-2), 1700111.

Ishibuchi, H., Tsukamoto, N., & Nojima, Y. (2016). Evolutionary many-objective optimization. *IEEE CEC*, 1-8.

Jin, W., Barzilay, R., & Jaakkola, T. (2018). Junction tree VAE for molecular graph generation. *ICML*, 2323-2332.

Kadurin, A., Nikolenko, S., Khrabrov, K., Aliper, A., & Zhavoronkov, A. (2017). druGAN: GAN-based de novo design. *Molecular Pharmaceutics*, 14(9), 3098-3104.

Kingma, D. P., & Welling, M. (2014). Auto-encoding variational Bayes. *arXiv*.

Li, H., & Zhang, Q. (2009). Multiobjective optimization: MOEA/D and NSGA-II. *IEEE Transactions on Evolutionary Computation*, 13(2), 284-302.

Nigam, A., Friederich, P., Krenn, M., & Aspuru-Guzik, A. (2020). Augmenting genetic algorithms with neural networks. *arXiv*.

Nigam, A., Polykovskiy, D., & Aspuru-Guzik, A. (2021). Janus molecular generation framework. *Machine Learning: Science and Technology*, 2(3), 035021.

Nigam, A., Polykovskiy, D., Kamya, P., Bjerrum, E. J., Engkvist, O., & Chen, H. (2022). Janus: Parallel tempered GA for molecular design. *Journal of Chemical Information and Modeling*, 62(9), 2064-2076.

Olivecrona, M., Blaschke, T., Engkvist, O., & Chen, H. (2017). Molecular de novo design via reinforcement learning. *Journal of Cheminformatics*, 9(1), 48.

Oulladji, L., Saadallah, M., Guellil, Z., Abbad, H., & Bekhti, H. (2025). AI-driven molecule generation and bioactivity prediction: A multi-model approach. *Journal of Computational Chemistry*, 46(8), 512–525.

Paul, S. M., Mytelka, D. S., Dunwiddie, C. T., Persinger, C. C., Munos, B. H., Lindborg, S. R., & Schacht, A. L. (2010). Improving R&D productivity. *Nature Reviews Drug Discovery*, 9(3), 203-214.

Polishchuk, P. G., Madzhidov, T. I., & Varnek, A. (2013). Size of drug-like chemical space. *Journal of Computer-Aided Molecular Design*, 27(8), 675-679.

Schneider, G. (2018). Automating drug discovery. *Nature Reviews Drug Discovery*, 17(2), 97-113.

Schneider, P., Walters, W. P., Plowright, A. T., Sieroka, N., Listgarten, J., Goodnow, R. A., ... Schneider, G. (2020). Rethinking drug design in the AI era. *Nature Reviews Drug Discovery*, 19(5), 353-364.

Segler, M. H. S., Preuss, M., & Waller, M. P. (2018). Planning syntheses with neural networks. *Nature*, 555(7698), 604-610.

Stahl, F., Mendez-Lucio, O., Vuckovic, D., Zdrazil, B., & Engkvist, O. (2022). Hybrid GA approaches for drug design. *Drug Discovery Today*, 27(8), 2203-2213.

Vamathevan, J., Clark, D., Czodrowski, P., Dunham, I., Ferran, E., Lee, G., ... Bender, A. (2019). Machine learning in drug discovery. *Nature Reviews Drug Discovery*, 18(6), 463-477.

Winter, R., Montanari, F., Noé, F., & Clevert, D. A. (2019). Learning continuous molecular descriptors. *Chemical Science*, 10(6), 1692-1701.

Yoshimori, A., Takeda, S., & Oono, K. (2021). Chemically valid GA operators. *Journal of Cheminformatics*, 13(1), 34.

You, J., Liu, B., Ying, R., Pande, V., & Leskovec, J. (2018). Graph convolutional policy networks. *NeurIPS*.

Zhang, Q., & Li, H. (2007). MOEA/D. *IEEE Transactions on Evolutionary Computation*, 11(6), 712-733.

Author(s) Information

Abbad Houda

Djillali Liabes University of Sidi Bel Abbes,
Computer Science Department
EEDIS Laboratory, Algeria
Contact e-mail: *houda.abbad@univ-sba.dz*

Taieb Brahim Mohammed

Djillali Liabes University of Sidi Bel Abbes,
Computer Science Department
EEDIS Laboratory, Algeria

Oulladji Latefa

Djillali Liabes University of Sidi Bel Abbes,
Computer Science Department
EEDIS Laboratory, Algeria

Abbad Leila

LIT Laboratory
ENSTA High School, Algiers; Algeria

To cite this article:

Houda, A., Mohammed, T.B., Latefa, O., & Leila, A. (2025). AI-driven molecular design: Synergizing deep generative models with evolutionary optimization. *The Eurasia Proceedings of Science, Technology, Engineering and Mathematics (EPSTEM)*, 38, 512-524.