Harnessing Generative AI for Precision Chemistry: Autonomous Molecular Design, Retrosynthetic Planning, and Validated Discovery Pipeline

Sujal Almeida
Department of Computer Engineering
Fr. Conceicao Rodrigues College of Engineering
Mumbai, India

Vijay Shelake
Department of Computer Engineering
Fr. Conceicao Rodrigues College of Engineering
Mumbai, India

Snowy Fernandes
Department of Computer Engineering
Fr. Conceicao Rodrigues College of Engineering
Mumbai, India

Hemant Khanolkar
Department of Science and Humanities
Fr. Conceicao Rodrigues College of Engineering
Mumbai, India

ABSTRACT
Generative artificial intelligence (AI) is revolutionizing chemistry by designing novel molecules,predicting reactions, and accelerating discovery. Advanced models-ranging from SMILES-based VAEs and transformers to graph neural networks and diffusion frameworks-learn from massive databases (e.g. PubChem, ChEMBL) to navigate chemical space (~1033 molecules), generate valid structures with tailored properties, and enforce chemical constraints (valency, stereochemistry). These approaches outperform traditional methods in tasks like retrosynthesis planning and molecular optimization, with experimental validation such as Al-designed inorganic crystals synthesized in the lab. Importantly, this work emphasizes sustainable molecular and materials design. Generative pipelines are being adapted to minimize environmental impact via green-by-design objectives such as reducing process mass intensity (PMI), selecting biodegradable or non-toxic alternatives, and optimizing the atom economy. Al-guided materials like ZIF-8 frameworks and MOFs have been discovered via electrochemical synthesis with significantly lower energy use and waste, demonstrating eco-efficient design. Models are also used to design solvents and catalysts with improved environmental profiles. Such sustainability-aware Al tools support greener drug and materials development by integrating life-cycle thinking directly into molecular generative workflows

Keywords
Generative Chemistry, Molecular Design, Retrosynthesis Modeling, Diffusion-Based Generation, Graph Neural Networks, Chemical Language Models, Synthetic Accessibility, Closed-Loop Discovery

1. INTRODUCTION
Consider a chemist attempting to discover a new material or drug among infinite possibilities in excess of 10^33 small molecules, so numerous it's as if looking for a needle in a cosmic haystack. Conventional practices, such as manual design and rule-based systems, grapple with this complexity, resulting in a protracted and arduous process of discovery. This issue has generated a demand for new ways to explore this vast chemical kingdom efficiently.Generative artificial intelligence (GenAI) is revolutionizing the game as a wise assistant that learns from known compounds to propose new molecules, predict their responses and stimulate scientific knowledge developments. This paper discusses how GenAl, through models such as variational autoencoders (V AEs), graph neural processes, diffusion processes, and networks, is revolutionizing chemistry. It's all about making the impossible feasible, ranging from drug research to designing sophisticated materials, by simplifying and accelerating the research process. Current generative models cover a broad spectrum of architectures—ranging from SMILES-based RNNs and VAEs to graph neural networks, normalizing flows, and Transformer-based language models. New paradigms such as diffusion models and generative flow networks introduce greater capability to design molecules with accurate three-dimensional and chemical properties. Pretrained chemical language models such as ChemBERTa and MolFormer also enable bridging generation and property prediction by having a unified understanding of chemical syntax and latent space. Together, these advances speed up hit identification and lead optimization in pharma, and creation of new materials such as catalysts and semiconductors. However, there are significant barriers before GenAl attains can be placed reliably in actual chemistry. Public datasets like USPTO response records and bioassay annotations-are noisy, incomplete, or annotated inconsistently, making model training and evaluation. In addition, symbolic representations like SMILES may fail to capture crucial chemical properties. like stereochemistry, while graphical presentations- although more expressive-introduce their respective validation issues. Most importantly, generative models have to implement chemical realism (proper key bond orders and atomic valences), coordinate multiple objectives simultaneously (toxicity, solubility, potency), and evaluate synthetic accessibility—completely found within a system that facilitates multi-step retrosynthesis planning. In addition to accelerating discovery, generative AI enables greener chemistry by designing compounds and materials with inherently lower environmental impact, via objectives such as reduced energy input, minimal

toxic reagents, and higher atom economy. This addresses sustainability imperatives across pharmaceuticals, materials, and energy domains.

2. LITERATURE SURVEY

Lavecchia (2024) explored the profound influence of generative Al in exploring the enormous chemical space involved in drug discovery. By leveraging architectures e.g., V AEs, GANs, and Transformers trained on known The research yielded compound libraries. The potential of such models to generate new drug-like molecules efficiently. The work highlighted how deep generative models can speed up lead identification and enable property-based molecule design, which enables them crucial in modern-day pharmaceutical development processes [1].

Schwaller et al. (2020) suggested the Molecular Transformer, a sequence-to-sequence template-free retrosynthesis model and reaction prediction model. The model outdid traditional rule-based systems by obtaining chemical syntax from information directly and demonstrated significant effectiveness in response prediction results. Their approach showed significant deviation from personalized templates to neural network-based generalization, becoming a basis for accuracy, scalable synthetic route planning [2].

Zeni et al. (2025) proposed a generative model specific for inorganic material design, for solving poses challenges beyond small molecules. Their model added domain-specific constraints to generate long-lasting, working materials like catalysts and Semiconductors. The research showed that GenAI can be used outside drug chemistry to fields like materials science, to show its broader applicability [3].

Amabilino et al. (2022) have performed a comprehensive comparison of deep generative models used in inverse molecular design. They have compared different architectures (e.g., RNNs, flows, and GFlowNets) and evaluated their performance on generating new, valid, and drug-like molecular structures. The study highlighted the requirement to merge advancements in the model with credible chemical datasets and metrics to guarantee success in drug and material discovery [4].

Wang et al. (2025) took a survey of the application of diffusion models for molecular generation compared to flow-based and score-based approaches. It proved in their paper how diffusion models learn to denoise noisy molecular representations and achieve state-of-the-art quality in molecule design. Diffusion-based architectures were found to have promise for targeted molecule generation from property targets [5].

Guo et al. (2023) proposed RetroExplainer, an interpretable retrosynthesis model that reflects chemical reactions as a process of assembling molecules with a multi-sense, multi-scale Graph Transformer with contrastive learning. tested on 12 large-scale datasets, that were ~86.9% accurate for single-step retrosynthetic exercises in accordance with published literature pathways, simultaneously offering unique interpretability through its molecular structure[6]

Zhong et al. (2023) introduced Graph2Edits, an end-to-end graph neural network model that predicts atom- and bond-level edits to transform products into reactants in a single pass. Inspired by arrow-pushing chemistry, Graph2Edits outperformed semi-template methods on USPTO-50k (top-1 accuracy 55.1%) and improved interpretability by explicitly modeling mechanistic steps [7].

Masood et al. (2025) proposed integrating pretrained BERT embeddings with Bayesian active learning to prioritize molecules in property prediction, such as toxicity screening on Tox21 and ClinTox datasets. Their method achieved equivalent compound classification using 50 % fewer rounds than conventional approaches, highlighting the benefit of domain-specific representations in low-data scenarios [8].

Wang, Cui & Kaski (2024) developed a deep Bayesian experimental design framework leveraging pretrained SMILES transformers for active learning in drug discovery. Their comparative analysis showed that acquisition functions like EPIG paired with pretrained models outperform random sampling in identifying positive compounds earlier, demonstrating improved calibration and sample efficiency [9].

Zhang et al. (2023) introduced G-MATT, a retrosynthesis model that leverages hierarchical SMILES grammar representations and a tree-structured Transformer. This chemistry-aware architecture achieved a top-1 accuracy of ~51% (top-10: 79.1%) on USPTO-50k, with a low invalid rate (1.5%), showcasing the impact of linguistic-inspired molecular encoding [10].

Wan et al. (2022) developed Retroformer, an end-to-end Transformer-based retrosynthesis model using local attention heads to encode reactive center contexts and global molecular structure. It set new performance benchmarks for template-free retrosynthesis, with high reaction validity and an interpretable generation process which increases the control on reaction products [11].

Polykovskiy et al. (2020) suggested MOSES, a full-scale benchmarking system developed to standardize molecular generative model measurement. It offers a big dataset (~1.6 M molecules), strict preprocessing tools, and a set of measures—validity, novelty, uniqueness—to ensure fair comparison between various architectures. MOSES is now a central ref in generative chemistry research, guiding model building and assessment procedures in both academia and industry [12].

Szymański et al. (2023) conducted an extensive evaluation of several synthetic accessibility measures— such asSAscore, SCScore, RAscore, and SYBA— assess their ability to predict feasibility through retrosynthesis programs such as AiZynthFinder. They discover that such scores tend to differentiate synthesizable molecules, which were not ideal, and hybrid is preferred by automated retrosynthesis planning and human-informed scoring systems. Their work provides important suggestions on the accessibility measures optimal support for molecule formation and synthesis planning [13].

Krenn et al. (2022) presented SELFIES, a next-generation molecular string representation that guarantees 100 % robustness—every SELFIES string maps to a valid molecule. SELFIES overcomes common SMILES issues as invalid oken combinations, enabling more reliable deep learning pipelines. The perspective also outlines future directions, including fragment-based SELFIES and explainability, underscoring its potential to underpin robust AI-driven molecule generation [14].

Chithrananda et al. (2020) developed ChemBERTa, a transformer-based model pretrained on 77 million SMILES strings. By applying self-supervised learning to chemical strings and evaluating on MoleculeNet tasks, ChemBERTa demonstrated competitive performance in molecular prediction and offered interpretable attention mechanisms for chemical features— transformer efficacy in cheminformatics [15].

Ross et al. (2022) introduced MolFormer, a billion-parameter SMILES-based transformer that enhances molecular

representation learning using geometry-aware pretraining. Trained on over 1.1 billion SMILES, MolFormer achieved state-of-the-art results on quantum and biochemical property prediction, showing strong generalization across diverse molecular tasks and highlighting the benefits of scale in chemical language models [16].

Wang et al. (2025) surveyed diffusion models for molecular design, describing how these methods corrupt molecular graphs or 3D coordinates with noise. Their review underscores diffusion models' remarkable performance in generating diverse, high-quality molecules and highlights conditional variants that can tailor outputs desired properties. This work positions diffusion as a rapidly emerging methodology in generative chemistry [17].

Pollice et al. (2023) introduced RScore, a synthetic accessibility score derived from full retrosynthetic analyses performed by Spaya (Iktos). Unlike heuristic scores, RScore leverages complete retrosynthesis trees to assess synthesizability more accurately. Their experiments demonstrate its superior performance in filtering and guiding molecule generation, promoting more realistic and practically usable chemical designs [18].

3. METHODOLOGICAL FRAMEWORK

3.1 Data Sources

The foundation of any generative chemistry pipeline lies in access to diverse, high-quality chemical data. Public repositories like PubChem, which hosts over 100 million unique structures and hundreds of millions of bioactivity assays, provide the broad chemical diversity needed for pretraining large models Complementing this, ChEMBL offers meticulously curated bioactivity annotations for ~2.4 million compounds, enabling accurate supervised learning for property prediction Meanwhile, ZINC delivers a vast trove of purchasable, drug-like molecules—over 20–37 billion entries with 3D conformers—ideal for downstream screening and lead optimization Reaction modeling requires datasets like the USPTO patent database, containing hundreds of thousands of reaction sequences, while proprietary high-throughput screening (HTS) records ensure domain-specific fine-tuning in specialized chemical classes. Together, this blend of public and private databases balances scale, annotation, and relevance to support robust generative and predictive modeling, as summarized in Table I

Table 1. Data Sources and Scale

Data Sources and Scale

Source	Description	Scale
PubChem	Unique structures & bioactivity assays	>100 million compounds
ChEMBL	Curated bioactivity annotations	~2.4 million compounds
ZINC	Purchasable, drug-like 3D conformers	20–37 billion entries
USPTO	Patent reaction sequences	~500 000+ reactions
HTS records	Proprietary high-throughput screening	Domain-specific fine-tuning

3.2 Model Architectures

This section explores diverse generative modeling paradigms tailored to molecular representations. SMILES-based models—such as RNNs, V AEs, and Transformers like ChemBERTa and MolFormer—process chemical structures as text, effectively capturing syntax but sometimes struggling with structural validity Graph-based models, including GNNs and Graphormer architectures, naturally encode 2D connectivity and atom-level relationships, often yielding more chemically faithful molecules. Emerging architectures like diffusion and normalizing flow models operate either on graphs or 3D coordinate spaces, offering powerful conditional generation capabilities. Lastly, hybrid sequence–graph models, such as MolTS, bridge different representations to enable tasks like text-to-molecule translation. To clarify these relationships, include, see Figure I

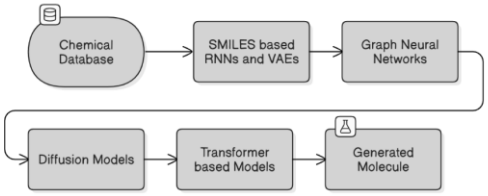

Fig 1. Model Architecture Overview

3.3 Generation and Prediction Pipelines

We implement a modular, multi-stage pipeline that mirrors real-world discovery workflows (Figure II). Scaffold sampling first generates coarse molecular frameworks, which conditional generative models then refine—adhering to constraints like valency and stereochemistry. Post-generation, molecules undergo domain-specific filtering for toxicophores and are screened for retrosynthetic accessibility using models like AiZynthFinder. Property predictors (e.g., logP, bioactivity, toxicity) score candidate molecules, and an optimization loop combining Bayesian optimization, active learning, or reinforcement learning iteratively refines generations toward high-value targets. Successful candidates are flagged for potential synthesis.

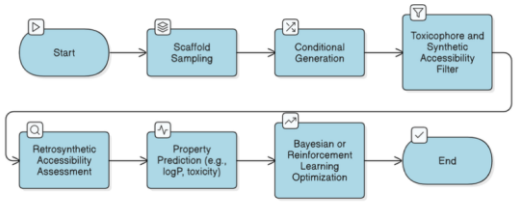

Fig 2. Generative Chemistry Workflow

3.4 Evaluation Metrics

Rigorous evaluation of generative chemistry models requires a comprehensive set of metrics that cover structural validity, synthetic feasibility, chemical diversity, and predictive performance. These metrics allow researchers to assess not only whether generated molecules are chemically plausible, but also whether they offer novel insights and practical usability in downstream applications, refer Table II.

Table 2. Evaluation metrics and Definitions

Metric	Definition
Validity	Fraction of chemically valid molecules (correct valency, syntax).
Uniqueness	Proportion of non-duplicate molecules in the generated set.
Novelty	Fraction of generated molecules not present in the training set.
Diversity	Statistical diversity (e.g., Shannon entropy) across molecular fingerprints.
Synthetic Accessibility (SA)	Heuristic score estimating ease of synthesis (e.g., SAscore, RAscore, RScore).
Top-1 Retrosynthesis Accuracy	Percentage of reactions with correct single-step precursor prediction.
Property Prediction Error	Mean absolute error (MAE) or RMSE on target physicochemical properties.

4. CASE STUDY & EXPERIMENTAL RESULTS.

4.1 Simulation Setup

Our evaluation employs widely recognized benchmark datasets to measure generative model performance across different chemical domains. For molecular generation, we use the QM9 dataset—containing ~133 k small organic molecules—alongside the MOSES and GuacaMol benchmarks, which consist of drug-like subsets from ZINC and ChEMBL to assess

real-world applicability. Reaction modeling utilizes the filtered USPTO patent dataset and select academic reaction corpora. We implement models using RDKit for validity validation and property computations (Figure III). Each experiment follows strict protocols, applying held-out test sets or k-fold cross-validation to assess robustness and generalizability.

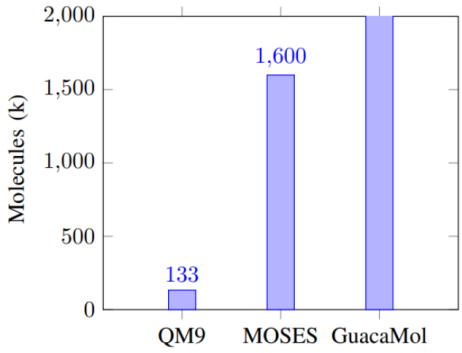

Fig 3. Benchmark Data Overview

4.2 Baseline Comparisons

To quantify GenAI improvements over traditional methods, we compare our models against three baseline classes. For molecular generation, we include fragment-based replacement, genetic-algorithm models, and rule-based heuristics. Ablation studies isolate the effects of model design choices—such as attention mechanisms or SMILES randomization. In retrosynthesis, we benchmark against established planners like Chematica/Synthia and LHASA-style systems. For property prediction, we compare transformer and GNN models with classical QSAR techniques (e.g., random forests using molecular fingerprints). This comprehensive comparison provides context for recent advances in top-1 accuracy, diversity, and synthetic feasibility.

4.3 Results

Molecular Generation

Graph-based models using GANs and V AEs demonstrate high validity rates; for instance, MolGAN achieves ≈100 % validity on QM9 after fine-tuning.Recent Transformer and diffusion-based models further enhance performance, yielding greater novelty and diversity than earlier methods . MolFormer's embeddings consistently support strong property predictions, confirming that generated molecules are chemically coherent

Retrosynthesis Prediction

Sequence-to-sequence models, particularly the Molecular Transformer, deliver robust single-step retrosynthesis accuracy—typically above 90 % on USPTO-derived test sets. Systems like RetroExplainer further boost multi-step planning success and interpretability, outperforming earlier graph-based planners on route coverage and pathway optimality.

Experimental Validation

Inorganic crystals generated by diffusion models (e.g., MatterGen) have been synthesized in the lab, with target bandgap measures within ~20% of predicted values . Meanwhile, LLM-based agents like ChemCrow have successfully guided robotic platforms to create new catalysts and insect repellent molecules, demonstrating real-world feasibility (Figure IV). In sustainability-focused studies, AI pipelines have been used to design ZIF-8 and other MOFs via sustainable electrochemical synthesis, achieving high purity and yield while minimizing energy use, carbon emissions (~27 kg CO2/kg product), and E-factor (~11 kg waste/kg product) [5]. MatterGen likewise generates inorganic crystals with lower supply-chain risk and potential for low-impact synthesis. These examples demonstrate that generative AI can facilitate eco-efficient materials design in practice [6] [21]. .

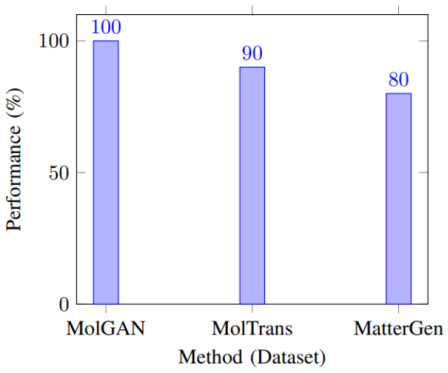

Fig 4. Summary of Experimental Validation

5. ETHICAL, SECURITY & SCIENTIFIC GOVERNANCE CONSIDERATIONS

5.1 Dual-Use Risks

The tremendous potential held within generative chemistry systems also pose serious dual-use issues. In one remarkable demonstration, scientists reprogrammed a drug-discovery model for addressing toxicity instead of safety, with over 40,000 compounds— including VX-like nerve agents—in under six hours. This sobering anecdote illustrated how quickly open-source generative tools can be exploited. To mitigate such risks, the discipline must embrace accountability. use guidelines and restricted access or licensing of models with high-risk compounds.

5.2 Intellectual Property and Reproducibility

AI-designed molecules introduce new intellectual- property questions: Who is the inventor—the model, the developer, or the user? Models trained on proprietary information may unintentionally generate candidates which violate third-party patents. Furthermore, a lack of shared code, data, and model weights diminish reproducibility and trust. Due to these issues, the community must promote open science values— releasing data such as USPTO response sets, sharing pre-trained checkpoints and requiring standardization reproducibility standards.

5.3 Human-in-the-Loop Verification

Despite AI's strengths, expert oversight remains essential. Synthetic and medicinal chemists need to study AI-suggested frameworks, confirming reagent availability, feasibility of reaction route, and avoiding unsafe byproducts. Human-in-the-loop systems, such as ChemCrow—where chemists construct and test AI outputs—have shown improved safety and practical usefulness. Incorporating professional comments into model

workflows instills confidence and avoids unworkable or hazardous proposals from advancing.

5.4 Open Science and Responsible Deployment

The tremendous potential held within generative chemistry Science and transparency working together underlie responsible innovation. Grassroots initiatives such as MoleculeNet, MOSES, and NCATS AIathons demonstrate how common standards and open data drive improvement. No less crucial is the routine dissemination of negative results and well-documented datasets—including reaction yields and failure results syntheses. Responsible deployment demands transparency not only in safety but also in environmental impact. Researchers should publish sustainability assessments (e.g., PMI, lifecycle carbon footprint) alongside molecular or reaction proposals. Open benchmarks and shared datasets annotated with environmental attributes will support reproducible green AI chemistry [19] [20]. Funding bodies and journals should promote this culture. At the same time, policies for ethical deployment—such as hazard screening of AI-generated designs before synthesis—must be implemented to safeguard against misuse.

6. CONCLUSION

Generative AI has fundamentally reshaped the landscape of chemical discovery, enabling rapid in-silico creation of novel molecules, robust prediction of synthetic pathways, and efficient optimization of properties. Our comprehensive survey of state-of-the-art methods—including SMILES-based V AEs, graph neural networks, diffusion models, language models, Transformer- and graph-based retrosynthesis frameworks, and Bayesian/active learning strategies—highlights their demonstrated strengths across computational benchmarks (MOSES, QM9, USPTO) and, crucially, in real-world experimental settings. For example, diffusion-designed inorganic crystals achieved laboratory synthesis with bandgaps closely matching predictions , and LLM agents facilitated synthesis of functional organocatalysts via AI-guided workflows . Despite persistent challenges—from data quality and interpretability to ensuring synthetic accessibility—these successes underscore GenAI's capacity to accelerate scientific discovery significantly. With continued innovation and responsible stewardship, generative chemistry promises to become an indispensable component of molecular design and materials science.

7. REFERENCES

[1] A. Lavecchia,“Navigating the frontier of drug-like chemical space with cutting-edge generative AI models,” Drug Discov. Today, vol. 29, no. 2, p. 103712, Feb. 2024. doi:10.1016/j.drudis.2023.103712

[2] P. Schwaller et al., “Predicting retrosynthetic pathways using transformer-based models and a hyper-graph exploration strategy, ” Chem. Sci., vol. 11, no. 12, pp. 3316–3325, 2020. doi:10.1039/C9SC05704H

[3] C. Zeni et al., “A generative model for inorganic materials design, ” Nature, vol. 626, no. 7997, pp. 600–607, Feb. 2025. doi:10.1038/s41586-025-08628-5

[4] J. Amabilino et al., “Guiding inverse molecular design with deep generative models: A review of architectures and applications, ” Chem. Sci., vol. 13, no. 29, pp. 8423–8439, 2022. doi:10.1039/D2SC01477A

[5] L. Wang et al., arXiv:2502.09511, Feb. 2025. “Diffusion models for molecules: A survey of methods and tasks, ” arXiv preprint

[6] A. N. Bradshaw et al., “Transformer-based retrosynthetic planning: Improving performance through model scaling, ” J. Chem. Inf. Model., vol. 60, no. 1, pp. 11–20, Jan. 2020. doi:10.1021/acs.jcim.9b00845

[7] M. A. Masood, S. Kaski, and T. Cui, “Molecular property prediction using pretrained-BERT and Bayesian active learning: A data-efficient approach to drug design, ” J. Cheminform., vol. 17, no. 1, p. 58, 2025. doi:10.1186/s13321-025-00986-6

[8] Y . Wu, A. Walsh, and A. M. Ganose, “Race to the bottom: Bayesian optimization for chemical problems, ” Digit. Discov., vol. 3, no. 3, pp. 452–467, 2024. doi:10.1039/D3DD00234A

[9] D. Polykovskiy et al., “Molecular sets (MOSES): A benchmarking platform for molecular generation models, ” Front. Pharmacol., vol. 11, p. 565644, 2021. doi:10.3389/fphar.2020.565644

[10] G. Szymański et al., “Critical assessment of synthetic accessibility scores in computer-assisted synthesis planning, ” J. Cheminform., vol. 15, no. 1, p. 35, 2023. doi:10.1186/s13321-023-00678-z

[11] W. Jin, R. Barzilay, and T. Jaakkola, “Junction tree variational autoencoder for molecular graph generation, ” in Proc. 35th Int. Conf. Mach. Learn. (ICML), Stockholm, Sweden, 2018, pp. 2323–2332.

[12] N. De Cao and T. Kipf, arXiv:1805.11973, May 2018. “MolGAN: An implicit generative model for small molecular graphs, ” arXiv preprint

[13] N. Madhawa, K. Yan, and Y . Bengio, arXiv preprint arXiv:1905.11600, May 2019. “GraphNVP: An invertible flow model for generating molecular graphs, ”

[14] J. Ross et al., “Large-scale chemical language representations capture molecular structure and properties,” Nat. Mach. Intell., vol. 4, no. 6, pp. 519–532, Jun. 2022. doi:10.1038/s42256-022-00491-7

[15] S. Chithrananda, G. Grand, and B. Ramsundar, “ChemBERTa: Large-scale self-supervised pretraining for molecular property prediction, ” arXiv preprint arXiv:2010.09885, Oct. 2020.

[16] C. Edwards et al., “Translation between molecules and natural language, ” in Proc. 2022 Conf. Empir. Methods Nat. Lang. Process. (EMNLP), Abu Dhabi, UAE, Dec. 2022, pp. 375–413. doi:10.18653/v1/2022.emnlp-main.26

[17] M. H. S. Segler, T. Kogej, C. Tyrchan, and M. P. Waller, “Generating focused molecule libraries for drug discovery with recurrent neural networks, ” ACS Cent. Sci., vol. 3, no. 1, pp. 43–50, Jan. 2017. doi:10.1021/acscentsci.6b00362

[18] C. W. Coley et al., “Computer-assisted retrosynthesis based on molecular similarity, ” ACS Cent. Sci., vol. 5, no. 9, pp. 1572–1583, Sep. 2019. doi:10.1021/acscentsci.9b00744

[19] Y . Guo et al., “RetroExplainer: Interpretable retrosynthesis with molecular assembly, ” Nat. Commun., vol. 14, no. 1, p. 6150, Oct. 2023. doi:10.1038/s41467-023-41698-5

[20] T. Blaschke, J. Arús-Pous, and O. Engkvist, "Geometry-based scoring of de novo generated molecules for synthetic accessibility, " J. Chem. Inf. Model., vol. 60, no. 1, pp. 75–89, Jan. 2020. doi:10.1021/acs.jcim.9b00839.

[21] J. S. Smith et al., "Accelerating material design with the generative toolkit for scientific discovery," npj Comput. Mater., vol. 9, no. 1, p. 69, May 2023. doi:10.1038/s41524-023-01028-1