Pharmacophore

ISSN-2229-5402

Journal home page: http://www.pharmacophorejournal.com

GENERATIVE TRANSFORMERS FOR CYP3A4-SPARING MOLECULE DESIGN USING ADMET AND SYNTHETIC FEASIBILITY CONSTRAINTS

Nguyen Thanh Huy[1]*, Pham Quang Minh[1], Le Thi Bich[2]

1. ***Department of Pharmaceutical Data Science, Faculty of Pharmacy, Vietnam National University, Hanoi, Vietnam.***
2. ***Department of Intelligent Drug Analytics, Faculty of Medicine, Can Tho University, Can Tho, Vietnam.***

ARTICLE INFO

Received:
20 November 2024
Received in revised form:
12 February 2025
Accepted:
15 February 2025
Available online:
28 February 2025

Keywords: Generative transformers, CYP3A4, ADMET, Molecular generation, Reinforcement learning, Synthetic feasibility

ABSTRACT

CYP3A4 is a major metabolic enzyme that can profoundly affect the oral exposure and clearance of drug candidates, making the design of molecules that avoid excessive CYP3A4 metabolism while retaining favorable drug-like properties a central challenge in early discovery. Traditional medicinal chemistry typically addresses metabolic liability only after candidate structures have been synthesized, a reactive approach that complicates the simultaneous optimization of CYP3A4 sparing, broader ADMET behavior, and synthetic feasibility. To address this, a conceptual generative transformer framework is proposed for designing novel CYP3A4-sparing molecules, capable of producing chemically valid candidates while respecting ADMET and synthetic constraints. The approach involves fine-tuning a transformer-based molecular language model trained on drug-like SMILES using reinforcement learning, with conditional generation guided by property tokens representing CYP3A4-sparing intent, ADMET desirability, and synthetic accessibility preferences. This framework could generate focused candidate libraries predicted to minimize CYP3A4 liability while maintaining favorable pharmacokinetic and synthetic characteristics, which would then be assessed for novelty, chemical diversity, and alignment with the desired design profile. By proposing metabolically resilient molecular starting points, a CYP3A4-sparing generative transformer could support hit-to-lead and lead-optimization decisions, complementing medicinal chemistry expertise without replacing the need for experimental validation.

To Cite This Article: Huy NT, Minh PQ, Bich LT. Generative Transformers for CYP3A4-Sparing Molecule Design Using ADMET and Synthetic Feasibility Constraints. Pharmacophore. 2025;16(1):31-9. https://doi.org/10.51847/Bl6094uJst

Introduction

CYP3A4 is a central enzyme in drug metabolism, and compounds with strong CYP3A4 liability may face unfavorable clearance, exposure variability, and interaction risk during development. In silico CYP modeling has therefore become an important component of early metabolic-risk assessment, with recent work addressing reversible inhibition, time-dependent inhibition, and machine-learning prediction of CYP3A4-related liabilities [1]. CYP-focused models such as CYPlebrity further illustrate how computational classifiers can support early recognition of cytochrome P450 inhibition patterns before extensive synthesis is undertaken [2]. For a CYP3A4-sparing design strategy, the economic and clinical motivation is to shift such assessment from late-stage triage toward the earliest stages of molecular ideation.

Current medicinal chemistry responses to CYP3A4 metabolism often rely on structural alerts, scaffold modification, matched molecular reasoning, and retrospective prediction after a molecule has already been proposed. Broader quantitative structure–activity relationship approaches for CYP inhibition show how computational alerts can inform prioritization, but they still depend on the quality and scope of the underlying training data [3]. ADMET platforms such as ADMETlab 2.0 offer integrated prediction environments for multiple pharmacokinetic and toxicity endpoints, yet these tools are typically used to evaluate existing structures rather than to generate new ones under explicit metabolic constraints [4]. This creates a gap between property prediction and proactive molecular invention.

Generative deep learning has changed this design logic by allowing models to construct molecules de novo rather than merely score fixed libraries. Variational autoencoders demonstrated that molecules could be embedded into continuous latent spaces

Corresponding Author: Nguyen Thanh Huy; Department of Pharmaceutical Data Science, Faculty of Pharmacy, Vietnam National University, Hanoi, Vietnam. E-mail: huy.nguyen@gmail.com.

for data-driven design [5], while conditional variational models showed that property goals could be incorporated directly into molecular generation [6]. Reinforcement-learning generators then made it possible to steer molecular distributions toward desired profiles by rewarding generated structures during sampling [7, 8]. Transformer-based chemical language models extend this idea by learning sequence dependencies in molecular strings and adapting them to property-guided generation [9, 10].

The thesis of this EAI article is that a decoder-only generative transformer could be conditioned to design molecules that are simultaneously CYP3A4-sparing, ADMET-aware, and synthetically feasible. The framework draws on molecular language modeling for SMILES generation [9], multi-objective reward shaping from reinforcement-learning molecular design [8], and modular synthetic feasibility assessment through reaction- and retrosynthesis-aware scoring [11-13]. Instead of treating CYP3A4 risk as a downstream filter, the proposed model embeds metabolic resilience into the generation objective itself. In this view, the transformer becomes a design engine that samples from a constrained chemical space shaped by metabolism, developability, and practical synthesis. To clarify the conceptual gap addressed in this article, **Table 1** summarizes how current CYP3A4-risk assessment strategies differ from the proposed EAI-guided generative design framework.

Table 1. Conceptual rationale for CYP3A4-sparing generative molecular design

Design challenge	Current approach	Limitation	Proposed EAI direction
CYP3A4 liability	In silico prediction of CYP3A4 inhibition, metabolism, and time-dependent inhibition	Often applied after molecules are already proposed or synthesized	Embed CYP3A4-sparing behavior directly into the molecular generation objective
Early metabolic-risk assessment	Structural alerts, QSAR models, and CYP-focused classifiers such as CYPlebrity	Dependent on training-data quality and may function mainly as retrospective filters	Use predictive CYP models as constraints or reward components during generation
ADMET optimization	Integrated platforms such as ADMETlab 2.0 evaluate pharmacokinetic and toxicity endpoints	Typically score existing structures rather than invent optimized candidates	Generate molecules under simultaneous CYP3A4, ADMET, and developability constraints
Molecular invention	VAEs, conditional generative models, reinforcement learning, and transformer-based SMILES models	Single-objective optimization can produce unrealistic or synthetically difficult molecules	Apply decoder-only transformers with multi-objective reward shaping
Synthetic feasibility	Retrosynthesis tools, reaction-aware scoring, and synthetic accessibility metrics	Feasibility is often assessed after generation	Include retrosynthesis- and reaction-aware scoring during molecule sampling
Overall design gap	Prediction and generation are often treated as separate stages	CYP3A4 risk remains a downstream triage problem	Build a proactive generative framework for CYP3A4-sparing, ADMET-aware, synthetically feasible molecules

Background

CYP3A4 Metabolism and Structural Avoidance Strategies

CYP3A4 is substrate-promiscuous, which makes its liabilities difficult to eliminate through a single universal rule. Computational models for CYP3A4 time-dependent inhibition highlight how structural patterns, chemical reactivity, and physicochemical context may all contribute to metabolic risk [1]. Practical avoidance strategies can include reducing exposed metabolic soft spots, adding steric shielding near labile atoms, modifying heterocycles, and balancing lipophilicity so that CYP recognition is less favorable. Because CYP3A4 liabilities are chemically diverse, a generative model should not merely exclude fixed alerts but should learn a flexible scoring landscape from curated CYP prediction models [2, 3].

Generative Transformers for Molecular Design

Transformer architectures adapt naturally to molecular design when SMILES strings are treated as chemical language sequences. MolGPT showed how a transformer-decoder model can generate molecules by learning next-token dependencies in SMILES notation [9], while chemical language models have shown that learned molecular syntax can support navigation through sparsely populated regions of chemical space [10]. Earlier recurrent and randomized-SMILES approaches established that molecular string generation benefits from robust sequence augmentation and learned chemical grammar [14, 15]. A CYP3A4-sparing transformer would build on these ideas by pre-training on broad chemical syntax and then conditioning generation toward metabolic and ADMET objectives.

Multi-Parameter Optimization in Drug Design

Drug design is inherently multi-objective because improving one property can degrade another, especially when potency, solubility, permeability, metabolic stability, and toxicity must be optimized together. Conditional molecular design methods show that property preferences can be encoded into generative models so that molecules are sampled from regions aligned with user-defined constraints [16]. Molecular property predictors such as directed message-passing models provide a way to score generated molecules against learned ADMET endpoints [17]. Integrated ADMET platforms and recent machine-learning ADMET systems further support the use of surrogate predictors as fast reward components for generative design [4, 18].

Synthetic Feasibility Scoring

Synthetic feasibility is essential because a generated molecule has limited value if it cannot be made through plausible chemistry. SCScore introduced a learned synthetic complexity measure derived from reaction data, allowing generated compounds to be penalized when they resemble unusually difficult synthetic targets [11]. Retrosynthetic planning systems such as AiZynthFinder provide a more route-aware perspective by connecting proposed structures to plausible disconnections and available precursors [12]. Retrosynthetic accessibility scoring further allows synthesis-informed feasibility to be approximated rapidly enough for generative loops [13].

Reinforcement Learning and Conditional Generation for Chemistry

Reinforcement learning allows a molecular generator to be updated according to a reward signal rather than only imitating a training distribution. Early de novo reinforcement-learning studies demonstrated that sequence-based molecular generators could be biased toward property objectives [7], and later deep reinforcement-learning approaches expanded this idea into multi-objective drug design [8]. REINVENT and its later developments show how policy-based generative chemistry can combine prior chemical knowledge with reward-directed sampling [19, 20]. Conditional generation and reinforcement learning are therefore complementary: tokens define the requested design region, while reward updates reshape the sampling policy toward molecules that satisfy it [16, 21].

Model Development Overview

High-Level Architecture

The proposed system uses a decoder-only transformer pre-trained on bioactive SMILES to learn molecular syntax, scaffold patterns, and local chemical dependencies. This backbone follows the molecular language-modeling logic of MolGPT [9] while drawing conceptually from broader chemical language models that support exploration of under-sampled chemical regions [10]. During fine-tuning, generated molecules would be scored by separate CYP3A4, ADMET, and synthetic feasibility modules, with reinforcement learning adjusting the generator toward higher-reward samples [8, 19]. Conditional property descriptors would be prepended to the input sequence so that the same generator could sample molecules under different metabolic and developability profiles.

Figure 1 illustrates the proposed CYP3A4-sparing generative transformer workflow, showing how molecular language pre-training, conditional property tokens, ADMET scoring, synthetic feasibility assessment, reinforcement learning, and medicinal chemistry review are integrated into a unified design pipeline.

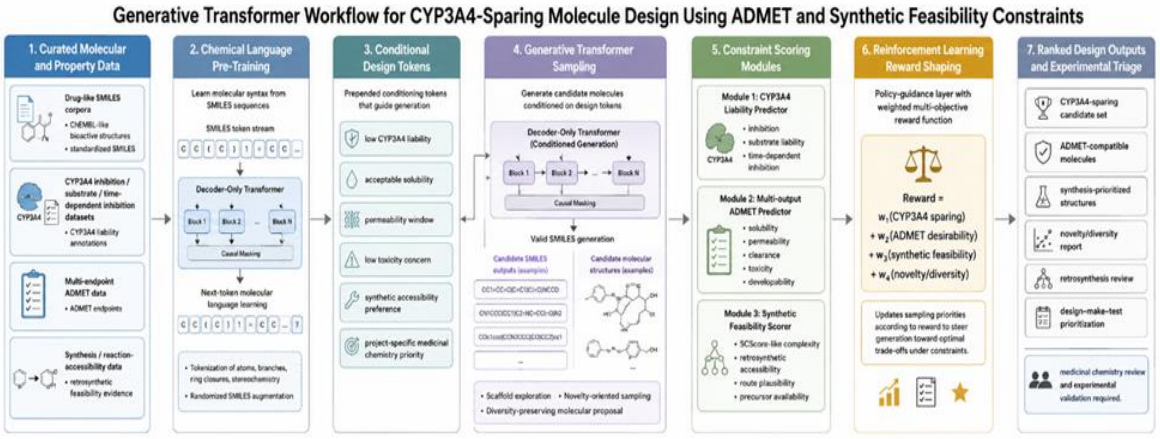

Figure 1. Generative Transformer Workflow for CYP3A4-Sparing Molecule Design Using ADMET and Synthetic Feasibility Constraints

Core Components

The framework contains a pre-trained SMILES generator, a CYP3A4 liability predictor, a multi-task ADMET predictor, a synthetic complexity scorer, and a reinforcement-learning controller. The CYP3A4 component would be informed by machine-learning models for CYP inhibition and time-dependent inhibition [1-3], while the ADMET component could follow multi-endpoint prediction strategies used in molecular representation learning and ADMET platforms [4, 17, 18]. Synthetic complexity would be represented using fast scoring functions such as SCScore and retrosynthesis-derived accessibility estimates [11, 13]. The controller would combine these outputs into a design reward, similar in spirit to reinforcement-learning de novo design systems [7, 8].

Table 2 defines the functional architecture of the proposed CYP3A4-sparing generative transformer, clarifying how each computational layer contributes to chemically valid, ADMET-aware, and synthetically feasible molecular design.

Table 2. Functional Architecture of the CYP3A4-Sparing Generative Transformer Framework

Framework layer	Primary function	Key inputs	Computational mechanism	Output generated	Analytical value for CYP3A4-sparing design
Molecular corpus preparation	Establishes the chemical language base for generation	Drug-like, lead-like, and fragment-like standardized SMILES; curated bioactive structures; randomized SMILES variants	Cleaning, canonicalization, tokenization, syntax-preserving augmentation	A chemically coherent sequence-learning corpus	Ensures that generated molecules remain syntactically valid and chemically plausible rather than arbitrary token strings
Decoder-only transformer backbone	Learns molecular syntax and scaffold-level sequence dependencies	Tokenized SMILES strings with atom, branch, ring, charge, and stereochemical symbols	Causal self-attention and next-token prediction	Pre-trained molecular language model	Provides the generative foundation for producing novel structures within learned medicinal chemistry space
Conditional property-token layer	Converts design intent into model-readable generation context	Tokens for low CYP3A4 liability, ADMET desirability, solubility, permeability, toxicity avoidance, and synthetic accessibility	Prepended control tokens that condition subsequent SMILES generation	Property-directed candidate sampling	Allows chemists to request different metabolic and developability profiles without rebuilding the model
CYP3A4 liability module	Estimates whether a generated molecule may present CYP3A4-related risk	Generated molecular structures; CYP3A4 inhibition, substrate, and time-dependent inhibition data	CYP-focused machine-learning classifier or predictor	CYP3A4-sparing probability or liability score	Moves CYP3A4 risk assessment from retrospective filtering into the generation objective
Multi-output ADMET module	Evaluates broader developability constraints	Molecular descriptors, graph features, or learned representations	Multi-task prediction across pharmacokinetic and toxicity endpoints	ADMET desirability profile	Prevents narrow CYP3A4 optimization from producing molecules with poor solubility, permeability, clearance, or toxicity profiles
Synthetic feasibility module	Determines whether generated molecules are practically reachable	Molecular structures, synthetic complexity estimates, retrosynthetic accessibility scores, route-planning evidence	SCScore-like complexity scoring, retrosynthesis-aware scoring, precursor and route plausibility checks	Synthetic feasibility score or triage category	Penalizes attractive but impractical molecules that lack plausible preparation routes
Reinforcement-learning controller	Aligns the generator with multi-objective medicinal chemistry priorities	CYP3A4 score, ADMET score, synthesis score, novelty, diversity, and validity metrics	Weighted reward shaping and policy-gradient fine-tuning	Updated generation policy	Balances metabolic sparing with developability and synthesis rather than optimizing a single property in isolation
Medicinal chemistry review layer	Provides expert interpretation and project-specific prioritization	Ranked molecules, predicted property profiles, novelty/diversity reports, retrosynthetic evidence	Human expert review and design–make–test triage	Experimentally testable candidate shortlist	Keeps the framework positioned as a hypothesis generator requiring expert judgment and laboratory validation

Design Principles

The first design principle is explicit conditioning, where property tokens narrow the generated chemical space before scoring occurs. The second is modularity, because the CYP3A4 model, ADMET panel, or synthetic scoring function could be updated independently as new assays or project-specific data become available, following the modular philosophy of modern generative platforms such as REINVENT 4 [20]. The third is multi-objective flexibility, where reward weights can represent different medicinal chemistry priorities without redefining the entire architecture. This modular strategy also allows retrosynthesis-aware filters from tools such as AiZynthFinder to be integrated after generation or within iterative scoring [12].

Data Sources and Molecular Representations

Pre-Training Corpus

The pre-training corpus would consist of cleaned, standardized, drug-like, lead-like, and fragment-like SMILES drawn from curated chemical databases. ChEMBL is a natural foundation because it organizes bioactivity data and chemical structures in a form widely used for machine-learning drug discovery [22]. Randomized SMILES augmentation could be used during pre-training to improve the robustness of sequence learning and reduce overdependence on a single canonical representation [15].

The resulting tokenizer would need to preserve chemically meaningful symbols, ring closures, stereochemistry, and branch structure so that the transformer learns valid molecular syntax rather than arbitrary character patterns [9].

CYP3A4 and ADMET Datasets
The CYP3A4 training data would include curated molecules annotated for inhibition, substrate behavior, or time-dependent inhibition, with careful attention to assay context and endpoint definition. CYPlebrity and related CYP modeling studies show how enzyme-specific inhibition data can be converted into machine-learning classifiers for early metabolic liability prediction [2, 3]. For broader ADMET modeling, descriptors or graph-based representations could be trained against multiple endpoints using approaches similar to molecular property prediction frameworks [17]. Public ADMET resources and platforms such as ADMETlab 2.0 and ADMET-AI illustrate how multiple pharmacokinetic and safety endpoints can be assembled into practical prediction systems [4, 18].

Synthetic Feasibility Data and Metrics
Synthetic feasibility data would combine calculated accessibility metrics, reaction-derived complexity scores, and retrosynthetic planning outputs. SCScore provides a learned estimate of synthetic complexity from reaction corpora, making it suitable as a rapid reward component in a generative loop [11]. Retrosynthesis-based systems such as AiZynthFinder offer complementary evidence by proposing plausible routes and connecting molecules to reaction knowledge [12]. RAscore provides a faster approximation of retrosynthetic accessibility, allowing a generator to penalize structures that are chemically attractive but route-infeasible [13].

Generative Transformer Architecture
Model Backbone and SMILES Tokenization
The model backbone would be a GPT-style decoder transformer trained with causal self-attention to predict the next token in a SMILES sequence. MolGPT provides the closest architectural precedent for this design, showing how transformer decoding can be applied directly to molecular generation [9]. Chemically aware tokenization would represent atoms, brackets, branches, ring indices, charges, and stereochemical symbols in ways that preserve the rules of molecular syntax. Randomized SMILES strategies could further help the model learn alternative string representations of the same molecule, improving generalization during sampling [15].

Conditional Generation with Property Tokens
Conditional generation would be implemented by prepending discrete property tokens that describe the desired design state, such as low CYP3A4 liability, acceptable solubility, or favorable synthetic accessibility. Conditional molecular generative models have shown that property information can guide sampling toward user-specified regions of chemical space [6, 16]. In the proposed transformer, these tokens would act as context for the entire generated sequence, shaping each subsequent SMILES token through self-attention. This design allows a single model to generate different molecular profiles by changing the conditioning string rather than retraining the backbone.

Reinforcement Learning Fine-Tuning Layer
After conditional pre-training, the transformer would be fine-tuned with a policy-gradient reinforcement-learning layer that rewards generated molecules satisfying the CYP3A4, ADMET, and synthesis objectives. Molecular de novo design through reinforcement learning provides the conceptual basis for adjusting a generator according to property rewards [7], while deep reinforcement-learning studies show how such rewards can be extended to multi-objective molecular optimization [8]. Memory-assisted and modern REINVENT-style frameworks further indicate how prior chemical knowledge can be retained while sampling is steered toward desirable chemical regions [19–21]. In this model, the reward would remain conceptual and project-adjustable, functioning as a guide for design exploration rather than as a claim of experimentally validated performance.

Incorporating ADMET and Synthetic Feasibility Constraints
Multi-Output ADMET Predictor
A multi-output ADMET predictor would serve as a rapid surrogate for estimating whether generated molecules fall within a desirable developability profile. Molecular representation studies show that learned features can support prediction across multiple property endpoints, making them suitable for integration into scoring pipelines [17]. ADMETlab 2.0 and ADMET-AI illustrate how diverse absorption, distribution, metabolism, excretion, and toxicity endpoints can be evaluated in unified computational frameworks [4, 18]. In the proposed model, these predictions would be converted into a desirability score that guides the transformer without implying that the generated structures have been experimentally validated.

Synthetic Feasibility as a Filter and Reward Modifier
Synthetic feasibility would be incorporated both as a hard filter and as a reward modifier to discourage chemically unrealistic structures. SCScore provides a reaction-informed estimate of synthetic complexity that can be evaluated quickly during generation [11]. Retrosynthesis-aware approaches such as AiZynthFinder and RAscore would add a route-based perspective,

helping identify molecules that are not only syntactically valid but also more plausible to make [12, 13]. In practice, the synthetic term would suppress structures that appear attractive in predicted ADMET space but lack a credible path to preparation.

Balancing CYP3A4 Sparing with Other Properties

Balancing CYP3A4 sparing with other properties requires a weighted reward rather than a single-objective filter. CYP3A4 predictors focused on inhibition and time-dependent liability can provide metabolism-specific signals [1–3], while ADMET prediction models can represent broader pharmacokinetic and safety preferences [4, 18]. The generative objective would allow medicinal chemists to adjust these weights according to project context, such as prioritizing metabolic stability in one series while emphasizing solubility or permeability in another. This flexibility is important because CYP3A4 avoidance should not create molecules that are otherwise unsuitable for development.

Reward Design and Conditional Generation Strategy

Multi-Objective Reward Shaping

The reward function would harmonize CYP3A4-sparing likelihood, ADMET desirability, and synthetic feasibility into a single policy-guidance signal. Reinforcement-learning molecular design has shown that generative models can be shifted toward user-defined objectives when rewards are assigned to generated structures [7, 8]. Additional molecular optimization strategies, including actor–critic approaches, support the broader idea that reward shaping can bias a generator toward preferred chemical regions without requiring explicit enumeration of all acceptable molecules [23]. In this framework, the reward should be smooth in expectation, chemically interpretable, and adjustable enough to reflect medicinal chemistry judgment.

Iterative Sampling and Model Update

During iterative sampling, the transformer would generate candidate SMILES, the scoring modules would assign reward components, and the policy would be updated to increase the probability of molecules aligned with the requested profile. REINVENT-style frameworks demonstrate how prior molecular knowledge can be combined with reward-driven exploration to avoid drifting into chemically implausible regions [19, 20]. Conditional transformer strategies and reaction-aware reinforcement learning further suggest that molecular generation can be guided by both desired properties and practical chemistry constraints [24]. This update cycle would be used conceptually to move the generator toward CYP3A4-sparing, ADMET-compatible, and synthetically feasible chemical space.

Integration into the Drug Design Pipeline

From Virtual Library to Design–Make–Test Cycle

The model would output a ranked virtual library rather than a final drug candidate, and medicinal chemists would review the generated structures for novelty, interpretability, and project fit. Benchmarking frameworks such as GuacaMol and MOSES emphasize the importance of assessing generated molecules through validity, novelty, diversity, and property alignment before prioritizing compounds for synthesis [25, 26]. Synthetic feasibility tools would then help triage the virtual library toward structures with plausible preparation routes [11-13]. The most compelling candidates could enter a design–make–test–analyze cycle where CYP3A4 liability and ADMET behavior are measured experimentally.

Feedback Loop for Model Refinement

Experimental results from synthesized compounds would provide feedback for retraining the CYP3A4 predictor, updating the ADMET surrogate, and refining the generative policy. This closed-loop logic is consistent with reinforcement-learning molecular design, where model behavior is progressively shaped by improved scoring information [8, 19]. Conditional generation would also allow newly learned project constraints to be expressed through updated property tokens rather than requiring a completely new generator [9, 16]. Over time, the system would be expected to align more closely with the medicinal chemistry series it supports, while still requiring experimental confirmation at each decision point.

Evaluation Strategy

Generation Quality Metrics

Generation quality would be assessed through chemical validity, uniqueness, novelty relative to training chemistry, and internal diversity. GuacaMol provides a framework for benchmarking de novo molecular design models using standardized quality and goal-directed tasks [25]. MOSES similarly supports evaluation of molecular generators by comparing generated distributions with reference chemical sets [26]. These metrics would establish whether the transformer is producing reasonable chemical structures before any claim is made about CYP3A4 sparing or ADMET utility.

Table 3 provides an evaluation and deployment-readiness framework for determining whether generated CYP3A4-sparing molecules are chemically valid, diverse, ADMET-compatible, synthetically feasible, and appropriate for experimental triage.

Table 3. Multi-Objective Evaluation and Deployment Readiness Framework for Generated CYP3A4-Sparing Molecules

Evaluation domain	Core question addressed	Recommended assessment criteria	Failure mode detected	Decision-use implication

Chemical validity	Are the generated outputs chemically interpretable structures?	Valid SMILES rate, valence correctness, stereochemical consistency, absence of chemically impossible fragments	Invalid strings, unstable structures, unrealistic atom connectivity	Invalid candidates should be removed before any ADMET, CYP3A4, or synthesis scoring is interpreted
Novelty and scaffold diversity	Does the model propose new but chemically meaningful candidates?	Novelty relative to training data, internal diversity, scaffold distribution, analog-series balance	Memorization of training molecules, excessive scaffold repetition, mode collapse	Diversity monitoring prevents the reward function from narrowing exploration to a small set of over-optimized chemotypes
CYP3A4-sparing performance	Do generated molecules appear less likely to show CYP3A4 liability?	Predicted inhibition risk, predicted substrate liability, time-dependent inhibition alerts, structural soft-spot analysis	False confidence from biased CYP predictors, reward hacking, hidden metabolic liabilities	CYP3A4-sparing scores should guide prioritization but must not replace experimental metabolism assays
ADMET compatibility	Are CYP3A4-sparing candidates still developable across other properties?	Solubility, permeability, clearance, toxicity risk, physicochemical balance, drug-likeness indicators	Molecules that avoid CYP3A4 but fail broader pharmacokinetic or safety expectations	ADMET compatibility prevents the model from solving metabolism at the expense of overall developability
Synthetic feasibility	Can the generated molecules plausibly be made?	Synthetic complexity score, retrosynthetic accessibility, route plausibility, precursor availability, reaction-step burden	High-scoring virtual molecules with unrealistic or unavailable synthetic routes	Feasibility scoring enables transition from virtual generation to design–make–test planning
Reward balance and interpretability	Is the model optimizing the intended medicinal chemistry profile rather than exploiting one metric?	Reward-component decomposition, weight-sensitivity analysis, Pareto-front inspection, property-trade-off mapping	Single-objective domination, artificial score inflation, loss of medicinal chemistry relevance	Transparent reward analysis helps chemists adjust design priorities according to project context
Prospective experimental readiness	Which generated molecules are suitable for laboratory follow-up?	Expert review, route confidence, structural novelty, ADMET/CYP profile alignment, assay prioritization	Overreliance on surrogate predictions without synthesis or testing	The framework should produce experimentally testable hypotheses, not unvalidated claims of drug-likeness
Model governance and updating	Can the system improve safely as new project data emerge?	Predictor calibration, dataset refresh, assay-context tracking, documentation of reward weights, version control	Stale predictors, dataset bias, non-reproducible generation settings	Governance ensures that generated candidates remain traceable, interpretable, and aligned with evolving project evidence

Constraint Satisfaction

Constraint satisfaction would examine whether generated molecules meet the intended CYP3A4-sparing, ADMET, and synthetic feasibility requirements according to the surrogate scoring modules. Conditional molecular design and MolGPT-style generation provide precedents for steering molecules toward specified property regions [9, 16]. Transformer-based chemical language models and adversarial latent-space generators also show how generated chemical spaces can be shaped by learned representations, although constraint satisfaction must still be verified with independent predictors and expert review [10, 27]. In this article, such evaluation would remain conceptual and comparative, avoiding unsupported claims about numerical success.

Prospective Synthesis and Testing

A prospective evaluation would select a small, chemically diverse set of generated molecules for synthesis and experimental profiling. Retrosynthetic planning and accessibility scoring would support this selection by prioritizing compounds with plausible routes before laboratory work begins [11, 12, 28]. Once synthesized, the molecules would be tested for CYP3A4 inhibition, metabolic stability, and relevant ADMET properties using assays appropriate to the discovery program [1-3, 29]. The purpose would be to evaluate whether the conceptual framework can generate experimentally useful hypotheses, not to assert performance before validation.

Limitations

Dependence on Surrogate Predictors

The quality of the generated design depends strongly on the reliability of the CYP3A4, ADMET, and synthetic feasibility predictors used in the reward function. CYP models can be sensitive to assay definitions, chemical series bias, and endpoint selection, which may lead the generator toward artifacts if the surrogate model is poorly calibrated [1-3, 30]. ADMET predictors face similar challenges because different endpoints may be measured under heterogeneous experimental conditions and may not transfer cleanly across chemical classes [4, 17, 18, 31]. For this reason, the proposed system should be treated as a hypothesis generator whose outputs require expert interpretation and experimental testing.

Chemical Space Coverage and Mode Collapse

A transformer generator may overrepresent familiar scaffolds from its pre-training distribution or collapse toward narrow chemotypes that exploit the reward function. Molecular generation benchmarks highlight the need to examine diversity and novelty, not only apparent property satisfaction [25, 26]. Memory-assisted reinforcement learning and modern generative platforms offer mechanisms for preserving prior chemical knowledge while encouraging exploration, but they do not remove the need for careful monitoring [20, 21, 32]. Diversity-promoting strategies, scaffold analysis, and human review would therefore be necessary safeguards in a CYP3A4-sparing design workflow. As a methodological safeguard, the proposed CYP3A4-sparing transformer workflow would require explicit monitoring for scaffold bias, reward exploitation, chemotype collapse, and loss of exploratory breadth, as summarized in **Table 4**.

Table 4. Safeguards against narrow or biased molecular generation

Risk in transformer generation	Why it matters	Safeguard
Overrepresentation of familiar scaffolds	The model may reproduce chemotypes already common in its pre-training data	Track scaffold frequency and compare generated molecules with training-set chemistry
Reward exploitation	Molecules may satisfy the scoring function without being genuinely useful	Use multi-metric evaluation and human medicinal-chemistry review
Mode collapse toward narrow chemotypes	Apparent optimization may reduce chemical diversity and novelty	Monitor diversity, novelty, and scaffold distribution during generation
Loss of balanced exploration	The model may prioritize high-scoring structures over broader chemical-space search	Apply diversity-promoting strategies and memory-assisted reinforcement learning

Conclusion

A generative transformer for CYP3A4-sparing molecule design would combine molecular language modeling, ADMET-aware scoring, and synthetic feasibility constraints in a unified design framework. Rather than treating CYP3A4 liability as a late-stage filter, the model would incorporate metabolic avoidance directly into the generation process. This shifts the role of computation from passive evaluation to proactive molecular proposal.

The main strength of the approach is its ability to generate candidate structures that are pre-shaped by multiple medicinal chemistry priorities. Conditional tokens would allow chemists to request different property profiles, while reinforcement learning would guide sampling toward molecules that better align with those profiles. Synthetic feasibility constraints would help ensure that proposed structures remain connected to practical chemical synthesis.

Important challenges remain before such a system could influence real discovery programs. Surrogate CYP3A4 and ADMET predictors must be carefully curated, calibrated, and updated with relevant experimental data. Synthetic route availability must also be assessed with realistic reaction knowledge rather than simple structural plausibility alone.

Open implementations, transparent benchmarks, and collaborative validation studies would help establish whether generative ADMET-aware design can become a dependable part of pharmaceutical research. A shared evaluation culture would also make it easier to compare transformer-based generators with other molecular design strategies. Ultimately, the value of this approach will depend on whether it can produce experimentally testable hypotheses that improve the efficiency and quality of early drug discovery.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Xu M, Lu Z, Wu Z, Gui M, Liu G, Tang Y, et al. Development of in silico models for predicting potential time-dependent inhibitors of cytochrome P450 3A4. Mol Pharm. 2022;20(1):194-205.
2. Plonka W, Stork C, Šícho M, Kirchmair J. CYPlebrity: Machine learning models for the prediction of inhibitors of cytochrome P450 enzymes. Bioorg Med Chem. 2021;46:116388.
3. Faramarzi S, Bassan A, Cross KP, Yang X, Myatt GJ, Volpe DA, et al. Novel (Q)SAR models for prediction of reversible and time-dependent inhibition of cytochrome P450 enzymes. Front Pharmacol. 2025;15:1451164.
4. Xiong G, Wu Z, Yi J, Fu L, Yang Z, Hsieh C, et al. ADMETlab 2.0: an integrated online platform for accurate and comprehensive predictions of ADMET properties. Nucleic Acids Res. 2021;49(W1):W5-14.
5. Gómez-Bombarelli R, Wei JN, Duvenaud D, Hernández-Lobato JM, Sánchez-Lengeling B, Sheberla D, et al. Automatic chemical design using a data-driven continuous representation of molecules. ACS Cent Sci. 2018;4(2):268-76.

6. Lim J, Ryu S, Kim JW, Kim WY. Molecular generative model based on conditional variational autoencoder for de novo molecular design. J Cheminform. 2018;10(1):31.
7. Olivecrona M, Blaschke T, Engkvist O, Chen H. Molecular de novo design through deep reinforcement learning. J Cheminform. 2017;9(1):48.
8. Popova M, Isayev O, Tropsha A. Deep reinforcement learning for de novo drug design. Sci Adv. 2018;4(7):eaap7885.
9. Bagal V, Aggarwal R, Vinod PK, Priyakumar UD. MolGPT: Molecular generation using a transformer-decoder model. J Chem Inf Model. 2022;62(9):2064-76.
10. Skinnider MA, Stacey RG, Wishart DS, Foster LJ. Chemical language models enable navigation in sparsely populated chemical space. Nat Mach Intell. 2021;3(9):759-70.
11. Coley CW, Rogers L, Green WH, Jensen KF. SCScore: synthetic complexity learned from a reaction corpus. J Chem Inf Model. 2018;58(2):252-61.
12. Genheden S, Thakkar A, Chadimová V, Reymond JL, Engkvist O, Bjerrum E. AiZynthFinder: a fast, robust and flexible open-source software for retrosynthetic planning. J Cheminform. 2020;12(1):70.
13. Thakkar A, Chadimová V, Bjerrum EJ, Engkvist O, Reymond JL. Retrosynthetic accessibility score (RAscore)–rapid machine learned synthesizability classification from AI driven retrosynthetic planning. Chem Sci. 2021;12(9):3339-49.
14. Segler MH, Kogej T, Tyrchan C, Waller MP. Generating focused molecule libraries for drug discovery with recurrent neural networks. ACS Cent Sci. 2018;4(1):120-31.
15. Arús-Pous J, Johansson SV, Prykhodko O, Bjerrum EJ, Tyrchan C, Reymond JL, et al. Randomized SMILES strings improve the quality of molecular generative models. J Cheminform. 2019;11(1):71.
16. Kang S, Cho K. Conditional molecular design with deep generative models. J Chem Inf Model. 2018;59(1):43-52.
17. Yang K, Swanson K, Jin W, Coley C, Eiden P, Gao H, et al. Analyzing learned molecular representations for property prediction. J Chem Inf Model. 2019;59(8):3370-88.
18. Swanson K, Walther P, Leitz J, Mukherjee S, Wu JC, Shivnaraine RV, et al. ADMET-AI: a machine learning ADMET platform for evaluation of large-scale chemical libraries. Bioinformatics. 2024;40(7):btae416.
19. Blaschke T, Arús-Pous J, Chen H, Margreitter C, Tyrchan C, Engkvist O, et al. REINVENT 2.0: an AI tool for de novo drug design. J Chem Inf Model. 2020;60(12):5918-22.
20. Loeffler HH, He J, Tibo A, Janet JP, Voronov A, Mervin LH, et al. Reinvent 4: Modern AI-driven generative molecule design. J Cheminform. 2024;16(1):20.
21. Blaschke T, Engkvist O, Bajorath J, Chen H. Memory-assisted reinforcement learning for diverse molecular de novo design. J Cheminform. 2020;12(1):68.
22. Mendez D, Gaulton A, Bento AP, Chambers J, De Veij M, Félix E, et al. ChEMBL: towards direct deposition of bioassay data. Nucleic Acids Res. 2019;47(D1):D930-40.
23. Wang Q, Wei Z, Hu X, Wang Z, Dong Y, Liu H. Molecular generation strategy and optimization based on A2C reinforcement learning in de novo drug design. Bioinformatics. 2023;39(11):btad693.
24. Nakamura S, Yasuo N, Sekijima M. Molecular optimization using a conditional transformer for reaction-aware compound exploration with reinforcement learning. Commun Chem. 2025;8(1):40.
25. Prykhodko O, Johansson SV, Kotsias PC, Arús-Pous J, Bjerrum EJ, Engkvist O, et al. A de novo molecular generation method using latent vector based generative adversarial network. J Cheminform. 2019;11(1):74.
26. Brown N, Fiscato M, Segler MH, Vaucher AC. GuacaMol: benchmarking models for de novo molecular design. J Chem Inf Model. 2019;59(3):1096-108.
27. Polykovskiy D, Zhebrak A, Sanchez-Lengeling B, Golovanov S, Tatanov O, Belyaev S, et al. Molecular sets (MOSES): a benchmarking platform for molecular generation models. Front Pharmacol. 2020;11:565644.
28. Honda S, Shi S, Ueda HR. SMILES Transformer: pre-trained molecular fingerprint for low data drug discovery. arXiv. 2019;1911.04738.
29. Maziarka Ł, Danel T, Mucha S, Rataj K, Tabor J, Jastrzębski S. Molecule Attention Transformer. arXiv. 2020;2002.08264.
30. Elton DC, Boukouvalas Z, Fuge MD, Chung PW. Deep learning for molecular design—a review of the state of the art. Mol Syst Des Eng. 2019;4(4):828-49.
31. Zhavoronkov A, Ivanenkov YA, Aliper A, Veselov MS, Aladinskiy AV, Aladinskaya AV, et al. Deep learning enables rapid identification of potent DDR1 kinase inhibitors. Nat Biotechnol. 2019;37(9):1038-40.
32. Sánchez-Lengeling B, Aspuru-Guzik A. Inverse molecular design using machine learning: Generative models for matter engineering. Science. 2018;361(6400):360-5.