Lead Optimization in Natural Product Chemical Space: A Methodological Model for Generative Design, Activity Prediction, Synthetic Feasibility, and Pharmacological Plausibility

Peter Nilsson[1]*, Eva Johansson[1], Lars Andersson[2]

[1]Department of AI for Cardioprotective Natural Products, Faculty of Pharmacy, Uppsala University, Uppsala, Sweden.
[2]Department of Computational Biophysics of Plant-Derived Peptides, Faculty of Pharmacy, Lund University, Lund, Sweden.

ABSTRACT

Lead optimization in natural product chemical space requires a methodological logic that is broader than molecular generation, predicted potency, or structural novelty. Natural products and natural product-inspired scaffolds can provide stereochemical richness, scaffold diversity, three-dimensionality, dense functionalization, and biosynthetically patterned motifs, but these same features complicate chemical representation, predictive reliability, synthetic planning, and pharmacological interpretation. Disconnected computational workflows may overvalue a generated molecule because it is chemically valid, predicted to be active, apparently novel, or heuristically accessible, even when its route feasibility, mechanism coherence, exposure plausibility, or uncertainty profile remains unresolved. This article proposes an original methodological model for lead optimization in natural product chemical space by integrating chemical-space representation, constrained generative design, activity prediction, uncertainty and applicability assessment, synthetic feasibility analysis, pharmacological plausibility evaluation, cross-domain conflict resolution, iterative redesign, and candidate-state assignment. The model treats generated structures as candidate hypotheses rather than optimized leads, interprets predicted activity relative to evidence credibility, distinguishes synthetic accessibility from practical synthetic feasibility, and evaluates pharmacological plausibility through target relevance, mechanism coherence, selectivity, exposure compatibility, polypharmacology, and safety-related concerns. The principal contribution is a qualitative decision architecture that prevents any single computational layer from dominating progression decisions and assigns candidates to transparent states: experimental prioritization, targeted optimization, feasibility-constrained redesign, evidence-insufficient hold, or rejection from the current optimization path. The model is proposed as a methodological framework and is not presented as prospectively validated.

Key Words: *Natural products, Lead optimization, Generative design, Activity prediction, Synthetic feasibility, Pharmacological plausibility*

eIJPPR 2025; 15(4):66-74

HOW TO CITE THIS ARTICLE: Nilsson P, Johansson E, Andersson L. Lead Optimization in Natural Product Chemical Space: A Methodological Model for Generative Design, Activity Prediction, Synthetic Feasibility, and Pharmacological Plausibility. Int J Pharm Phytopharmacol Res. 2025;15(4):66-74. https://doi.org/10.51847/ObdfX5Dk30

INTRODUCTION

Natural products continue to influence drug discovery because they provide biologically shaped structures that differ from many synthetic screening collections. Retrospective analysis of natural products has shown that they can contribute discovery-relevant patterns and structural diversity, while also requiring careful interpretation of historical success. Such evidence supports the view that natural product chemical space is important,

Corresponding author: Peter Nilsson
Address: Department of AI for Cardioprotective Natural Products, Faculty of Pharmacy, Uppsala University, Uppsala, Sweden.
E-mail: ✉ peter.nilsson@farmbio.uu.se
Received: 08 April 2025; **Revised:** 16 August 2025; **Accepted:** 22 August 2025

This is an **open access** journal, and articles are distributed under the terms of the Creative Commons Attribution-Non Commercial-ShareAlike 4.0 License, which allows others to remix, tweak, and build upon the work non-commercially, as long as appropriate credit is given and the new creations are licensed under the identical terms.

but not the view that natural origin alone establishes therapeutic value.

Lead optimization must therefore begin by treating natural products as structured opportunity spaces rather than automatically privileged candidates [1].

Recent drug-discovery literature has emphasized renewed opportunities for natural products through improved isolation, screening, informatics, biosynthetic analysis, and synthetic biology. These developments make natural products more compatible with contemporary discovery workflows, but they do not remove the need for rigorous candidate triage. A structurally attractive natural product-derived molecule may still fail because of poor exposure, weak selectivity, difficult synthesis, or uncertain mechanism. The methodological question is therefore how to convert natural product-inspired chemical opportunity into evidence-linked lead-optimization decisions [2].

Natural product-based drugs and chemical probes provide useful lessons for chemical-space reasoning because they show how natural product-derived structures can enter therapeutically or experimentally useful contexts. Cheminformatic analysis of such compounds indicates that relevant natural product-derived molecules may differ from generic small molecules in ways that matter for scaffold selection and optimization. However, evidence from successful drugs and probes should not be generalized to all natural product-like structures. The optimization process must preserve the distinction between observed utility in selected compounds and assumed utility in newly generated analogues [3].

Cheminformatics provides essential tools for representing natural product chemical space, including descriptors, similarity methods, scaffold analysis, and database-driven mapping. These methods can help convert complex natural product structures into computationally usable forms, but representation is not neutral. If stereochemical identity, ring complexity, biosynthetic motifs, or domain coverage are simplified too aggressively, downstream generation and prediction may operate on an incomplete version of the optimization problem. Natural product lead optimization therefore requires representation choices that are chemically meaningful and decision-relevant [4].

Natural product-inspired scaffold hopping illustrates both the promise and the risk of computational transformation. Similarity-based methods can connect natural products to synthetic mimetics and help explore analogues that retain selected structural or pharmacophoric features. Yet scaffold replacement does not prove activity preservation, and structural novelty does not establish lead quality. The purpose of the present article is to propose a methodological model that coordinates generation, activity prediction, synthetic feasibility, pharmacological plausibility, uncertainty, and redesign rather than treating any single layer as decisive [5].

Natural product chemical space

Natural product chemical space is not ordinary generic molecular space because its structures often reflect biosynthetic selection, stereochemical control, and dense functional organization. These properties can create unusual molecular shapes, interaction patterns, and scaffold types relevant to lead optimization. At the same time, such features may place candidates outside well-sampled regions of medicinal chemistry datasets. The model proposed here treats this dual nature as a central methodological condition: natural product complexity can be useful only when its consequences for prediction, synthesis, and pharmacology remain visible.

Scaffold diversity creates opportunities for discovering or modifying molecular frameworks that are underrepresented in conventional compound collections. However, scaffold diversity is not equivalent to pharmacological value. A scaffold may be novel, natural product-like, or structurally impressive while lacking target relevance, exposure compatibility, or practical synthetic access. For this reason, the model separates chemical-space admissibility from candidate prioritization.

Stereochemical richness and three-dimensionality are similarly double-edged. They may support selective recognition in biological systems, but they also increase the difficulty of descriptor design, conformational modeling, route planning, and experimental synthesis. A computational workflow that encodes a candidate only as a valid string or graph may fail to capture the stereochemical burden that determines whether an analogue can be made or meaningfully tested. Lead optimization in this space therefore requires structural representation that remains connected to practical chemistry.

Natural-product-likeness should be used as a contextual descriptor, not as a therapeutic claim. A candidate may be highly natural product-like because it shares ring systems, stereochemical density, or functional-group patterns with known natural products. This does not imply that the candidate has suitable activity, selectivity, pharmacokinetics, safety, or optimization potential. The proposed methodological model therefore uses natural-product-likeness to define admissible design regions, not to replace evidence of biological or developability plausibility.

Data sparsity and distribution shift are recurring challenges in natural product chemical-space modeling. Natural product-derived candidates may be distant from the molecules used to train activity-prediction, property-prediction, or retrosynthetic models. This distance can

make apparently precise computational outputs less reliable than their numerical form suggests. A chemical-space representation layer should therefore document not only what the molecule is, but also how well the relevant computational evidence is expected to apply to it.

Generative design and activity prediction

Generative design can support lead optimization by proposing analogues, scaffold variants, or transformed candidates under predefined objectives. Reinforcement-learning approaches show how molecular generation can be directed toward specified molecular goals. In a lead-optimization setting, however, this capability is useful only when the reward structure reflects biologically and chemically meaningful constraints. Otherwise, the model may optimize a surrogate objective while drifting away from plausible leads [6].

Latent-space design provides another route to molecular proposal by learning continuous representations that permit interpolation, optimization, and candidate sampling. Such methods can be powerful because they turn molecular design into a navigable computational problem. Yet latent-space proximity does not necessarily preserve mechanism, synthesizability, stereochemical realism, or assay relevance. In natural product chemical space, latent optimization must therefore be constrained by scaffold logic and chemical-space admissibility [7].

Prospective and target-focused examples of deep-learning-enabled molecular design show that generative workflows can contribute to inhibitor discovery in defined contexts. These examples are important because they demonstrate that computational design can participate in real discovery workflows when paired with experimental testing and target-specific evaluation. They should not be interpreted as evidence that any AI-generated molecule is automatically credible. The methodological model therefore treats generated structures as hypotheses that must pass additional gates [8].

Chemical validity is a necessary but insufficient condition for lead optimization. Robust molecular string representations can reduce invalid generated structures and improve the formal reliability of molecular encoding. This is valuable for candidate generation, but it does not address target engagement, feasibility of synthesis, exposure, selectivity, or toxicity. A generated molecule that is validly represented remains only an input to optimization, not an optimized lead [9].

Generative molecular design literature increasingly emphasizes the need for evaluation, constraints, and decision context beyond molecule production. Automated generative and reinforcement-learning workflows can propose candidate bioactive compounds, but their outputs still depend on training data, objective design, and downstream filters. In natural product lead optimization, generative design and activity prediction should therefore be coupled carefully rather than allowed to reinforce each other's errors. A candidate optimized toward predicted activity may become less realistic, less synthesizable, or less pharmacologically coherent [10].

Activity prediction adds another powerful but limited evidence layer. Molecular machine-learning benchmarks help compare property-prediction approaches, but benchmark performance should not be treated as proof of prospective lead quality [11, 12]. Bioactivity benchmarking has shown that deep learning can perform well under defined data conditions, yet such performance remains dependent on dataset composition and assay context [13]. Learned molecular representations can also influence prediction behavior, transferability, and endpoint performance, which means representation choice is part of the optimization logic rather than a neutral technical detail [14]. Graph-based property prediction and drug–target interaction modeling can support activity or target hypotheses, but these predictions must remain distinct from confirmed target engagement [15]. Drug–target interaction prediction may help prioritize mechanistic hypotheses, yet it still requires biological interpretation and validation [16]. Uncertainty-aware molecular property prediction provides a way to flag predictions that should not be treated as equally reliable across chemical space [17]. The major methodological opportunities and failure modes arising from the interaction of generative design and activity prediction are summarized in **Table 1**.

Table 1. Methodological Opportunities and Failure Modes in Generative Design and Activity Prediction for Natural Product Lead Optimization

Methodological component	Intended optimization function	Potential contribution	Major failure mode	Natural product-specific challenge	Evidence required before progression	Inappropriate inference to avoid
Constraint-aware molecule generation	Propose candidate analogues or scaffold variants within defined boundaries	Expands candidate options while preserving design intent	Reward hacking or objective overfitting	Complex scaffolds may be distorted by generic generation	Structural provenance, validity, admissibility, and constraint compliance	Generated structure equals optimized lead

Latent-space optimization	Explore learned molecular neighborhoods	Enables efficient candidate search	Latent proximity may not preserve mechanism or feasibility	Natural product motifs may be poorly represented	Chemical-space mapping and domain assessment	Latent optimization equals biological relevance
Sequence or string-based generation	Produce chemically valid molecular encodings	Improves formal candidate generation	Valid string may remain chemically or pharmacologically impractical	Stereochemistry and macrocyclic complexity may be underrepresented	Validity plus stereochemical and route review	Valid SMILES proves viability
Activity prediction	Estimate bioactivity or target-related signal	Adds prioritization evidence	Prediction outside applicability domain	Sparse natural product-like training coverage	Domain relevance, label quality, uncertainty, and assay context	Predicted potency proves activity
Drug–target interaction prediction	Suggest target-engagement hypotheses	Supports mechanistic prioritization	Target-context mismatch	Natural product polypharmacology may complicate interpretation	Orthogonal target evidence and biological plausibility	Predicted target interaction proves mechanism
Uncertainty assessment	Qualify prediction credibility	Prevents equal treatment of unequal predictions	Miscalibration or hidden distribution shift	Natural product candidates may be out-of-domain	Applicability analysis and uncertainty reporting	High confidence overrides domain mismatch
Multiobjective optimization	Balance activity, properties, feasibility, and plausibility	Prevents single-objective overoptimization	Unexplained composite score hides conflict	Structural richness may conflict with feasibility or exposure	Transparent trade-off rationale	Weighted score proves progression readiness

Synthetic feasibility and pharmacological plausibility

Synthetic feasibility must be evaluated independently from molecular generation and predicted activity. Computer-assisted retrosynthesis can support route reasoning by identifying possible disconnections and analogical synthesis logic. Such output is useful for early feasibility assessment, but it should not be treated as experimental proof that a molecule can be made. In natural product-derived optimization, this distinction is especially important because stereochemical control, dense functionality, and scaffold complexity can dominate practical execution [18].

Deep learning and symbolic AI have extended the scope of computational synthesis planning. These approaches can generate retrosynthetic routes and combine learned reaction knowledge with planning strategies. However, a route proposal does not automatically establish starting-material availability, yield reliability, chemoselectivity, purification feasibility, or scalability. The proposed model therefore interprets retrosynthesis as structured evidence requiring chemical review rather than as a binary pass [19]. Reaction prediction can further inform feasibility by estimating plausible transformations and, in some cases, uncertainty around predicted reaction outcomes. This helps connect virtual candidates to chemical precedent. Still, reaction prediction does not resolve the full burden of multi-step synthesis, protecting-group strategy, stereochemical control, or route robustness. Natural product lead optimization therefore requires route-level reasoning beyond isolated transformation plausibility [20]. Generative design can produce candidates that look attractive computationally but are difficult to synthesize. Analyses of molecules proposed by generative models show why synthesizability must be assessed as a separate methodological dimension. This is not merely a practical afterthought; it can determine whether an apparently promising candidate should progress, be redesigned, or be held. The model therefore places synthetic feasibility after generation and prediction but before final prioritization [21].

Synthetic accessibility and retrosynthetic accessibility scores are useful for triage, especially when large candidate sets must be reduced. Machine-learned synthesizability classification based on retrosynthetic planning can rapidly estimate whether candidates appear accessible. Yet such scores remain abstractions of practical chemistry and cannot substitute for route-level review. A favorable score should trigger deeper feasibility analysis rather than serve as proof of synthesis [22].

Pharmacological plausibility is broader than predicted potency and must be integrated with synthetic feasibility. ADMET prediction can help identify exposure, metabolism, and toxicity-related liabilities that affect whether a predicted activity signal is meaningful [23]. A candidate may be computationally active but synthetically implausible, synthetically feasible but pharmacologically

weak, structurally novel but mechanistically incoherent, or potent in prediction but exposure-incompatible. The methodological purpose is not to reject every imperfect candidate, but to determine whether the limitation is remediable through constrained redesign or whether it undermines the optimization path. The interacting evidence requirements that connect synthetic feasibility with pharmacological plausibility are organized in **Table 2**.

Table 2. Integrated Evidence Requirements for Synthetic Feasibility and Pharmacological Plausibility in Natural Product Lead Optimization

Evidence domain	Key methodological question	Relevant indicators	Major uncertainty	Cross-domain interaction	Progression implication	Inappropriate inference to avoid
Synthetic accessibility	Does the candidate appear chemically reachable?	Accessibility estimate, structural complexity, functional-group burden	Score may not reflect practical route risk	High activity may be unusable if accessibility is misleading	Use for triage, not final feasibility	Accessibility score proves synthesis
Retrosynthetic plausibility	Is there a credible route concept?	Disconnections, route precedent, starting materials, step logic	Generated route may fail experimentally	Route burden may force structural redesign	Proceed only with visible route risks	Route proposal proves laboratory success
Stereochemical control	Can required stereochemistry be established or preserved?	Chiral centers, stereoselective steps, conformational constraints	Stereochemical ambiguity may be underestimated	Natural product-like richness may increase route difficulty	Redesign if stereochemical burden is excessive	Correct drawing proves stereochemical feasibility
Chemoselectivity and protecting groups	Can functional groups survive the route?	Reactive motifs, protecting-group needs, incompatible transformations	Functional-group density may complicate execution	Feasibility may conflict with pharmacophore preservation	Modify structure only if mechanism rationale remains intact	Dense functionality is automatically beneficial
Target relevance	Is the predicted activity linked to a meaningful biological target or phenotype?	Disease context, target role, assay relevance	Target prediction may be context-dependent	Feasible synthesis is insufficient if target rationale is weak	Hold or redesign around stronger target logic	Predicted activity proves target engagement
Exposure and ADMET plausibility	Could the candidate reach a relevant biological context?	Solubility, permeability, metabolism, toxicity alerts	Predictions may be uncertain or endpoint-limited	Potency may be irrelevant if exposure is implausible	Advance only with testable ADMET assumptions	Potency alone establishes plausibility
Polypharmacology and safety	Are multi-target effects coherent or risky?	Selectivity, off-target signals, pathway consistency, toxicity concerns	Beneficial and harmful polypharmacology may overlap	Mechanistic promise may be offset by safety liabilities	Prioritize only if risks are visible and testable	Natural-product-like profile ensures safety

Proposed methodological model

The proposed methodological model begins with optimization-context definition rather than molecule generation. The target or phenotype, therapeutic objective, starting scaffold, desired property profile, known liabilities, and experimental constraints must be articulated before computational design begins. This sequence reflects the broader need for realistic decision framing in AI-enabled drug discovery, where computational tools should be connected to impact-relevant questions rather than used as autonomous discovery engines [24]. In this model, the optimization context defines what generation is allowed to attempt and what evidence will be required before progression.

The second layer translates natural product chemical space into explicit design constraints. These include scaffold retention or transformation boundaries, stereochemical requirements, novelty limits, target-relevant motifs, forbidden liabilities, and synthesis-aware restrictions. Chemical and biological data limitations must remain visible at this stage because the quality of subsequent predictions depends on what is represented and what is

missing [25]. The output is a constraint set that regulates generative design rather than allowing open-ended exploration.

Candidate generation is then used to propose analogues, scaffold variants, or constrained structures, not to declare success. Novelty is evaluated as a feature of the candidate's relationship to the design space, not as a substitute for biological or synthetic evidence. Pharmacological plausibility also enters early because natural product-inspired candidates may express multi-target behavior, and computational polypharmacology can support structured reasoning about beneficial or risky target profiles [26]. The model therefore treats generation as hypothesis production under chemical and biological constraints.

Activity prediction and affinity estimation are interpreted as evidence layers with defined limits. Ligand-based, graph-based, structure-based, and target-prediction approaches may support candidate comparison, but they do not prove target engagement. Molecular docking can contribute structural hypotheses, yet docking outputs should not be used as mechanistic confirmation without orthogonal evidence [27]. The model therefore separates predicted activity, predicted affinity, target-engagement hypotheses, and pharmacological plausibility.

The central architecture integrates six operations: chemical-space representation, constrained generation, activity prediction, uncertainty assessment, synthetic feasibility analysis, and pharmacological plausibility evaluation. Cross-domain integration then identifies conflicts rather than collapsing them into an unexplained composite score. For example, a candidate with strong predicted activity but weak route credibility is not equivalent to a candidate with moderate predicted activity, feasible synthesis, and coherent pharmacology. The model is intended to make such distinctions visible and actionable. The integrated methodological architecture and its alternative lead-optimization states are shown in **Figure 1**.

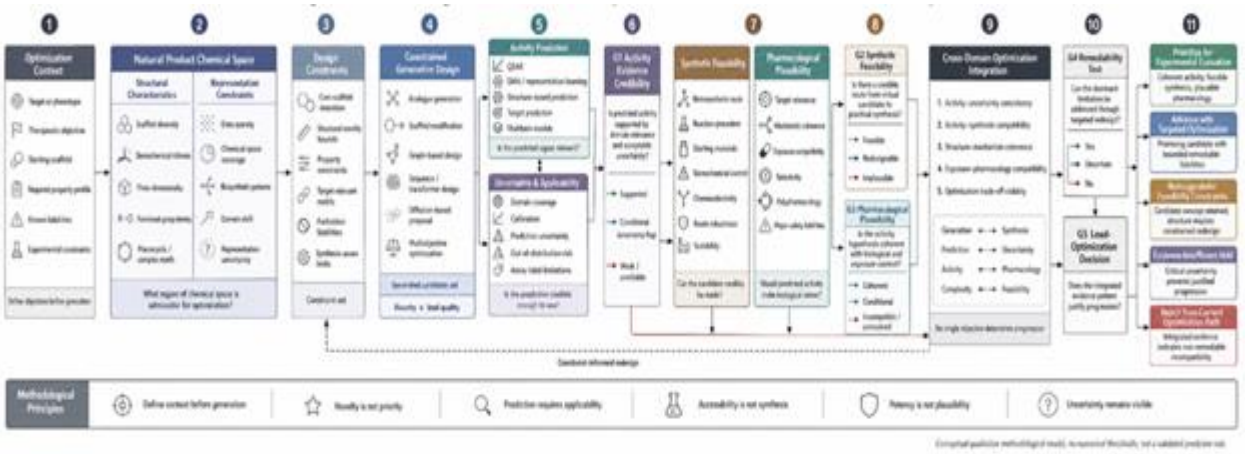

Figure 1. Integrated Methodological Model for Lead Optimization in Natural Product Chemical Space

Alt-text

A wide left-to-right methodological model begins with optimization objectives and natural product chemical-space representation. Candidates pass through constrained generative design, activity prediction, uncertainty evaluation, synthetic feasibility assessment, and pharmacological plausibility analysis. Cross-domain integration evaluates conflicts among activity, synthesizability, mechanism, exposure, and uncertainty. Explicit gates then direct candidates toward five alternative states: Prioritize for Experimental Evaluation, Advance with Targeted Optimization, Redesign Under Feasibility Constraints, Evidence-Insufficient Hold, or Reject from Current Optimization Path.

The model assigns candidates to exactly five alternative optimization states. "Prioritize for Experimental Evaluation" is reserved for candidates with coherent evidence across chemical-space admissibility, credible activity prediction, feasible synthesis, and plausible pharmacology. "Advance with Targeted Optimization" applies when the candidate remains promising but requires focused improvement of a bounded liability. "Redesign Under Feasibility Constraints" applies when the concept is valuable but the structure requires synthesis-aware redesign.

The remaining states prevent overinterpretation. "Evidence-Insufficient Hold" applies when uncertainty, assay relevance, target context, route credibility, or exposure plausibility is unresolved. "Reject from Current Optimization Path" applies when core assumptions fail or the incompatibility is not plausibly remediable under the current objective. These states are methodological categories for decision transparency, not empirically validated classifications. Their function is to guide what evidence should be generated next.

ISSN (Online) 2249-6084 (Print) 2250-1029

Validation considerations

Validation of the proposed model must be distinguished from validation of its individual components. Ligand-based classification benchmarks can reward memorization rather than generalization when related structures recur across training and test sets. This is particularly relevant for natural product-inspired optimization, where scaffold families, analogues, and close derivatives may create misleading estimates of predictive reliability. A credible evaluation strategy should therefore examine whether decisions remain stable under scaffold-aware and out-of-domain conditions [28].

Uncertainty-aware and conformal prediction approaches can help qualify whether model outputs are suitable for decision-making. Large-scale comparison of QSAR and conformal prediction methods shows the importance of applicability, confidence, and validation design in drug-discovery prediction workflows. However, uncertainty estimates should not be treated as universally reliable simply because they are numerically expressed. They require calibration, domain assessment, and interpretation relative to the specific chemical and biological context [29].

Retrospective validation can test whether the model's gate logic would have classified known candidate histories in a reasonable way. Such evaluation could examine whether successful candidates would have passed chemical-space, activity, feasibility, and plausibility gates, and whether failed candidates would have triggered hold, redesign, or rejection. This would not prove prospective utility, but it could reveal whether the decision architecture captures meaningful distinctions. Retrospective analysis should also include negative or discontinued candidates when available.

Component validation is different from candidate validation. An activity-prediction model may be validated on an external benchmark, a retrosynthesis tool may generate plausible routes, and an ADMET predictor may perform well on curated endpoints, yet a specific generated candidate may still fail experimentally. Candidate validation requires synthesis or acquisition, experimental activity testing, target-engagement evidence, and pharmacological evaluation appropriate to the intended mechanism. The proposed model therefore treats component-level evidence as necessary but not sufficient.

Prospective utility validation would require applying the methodological model to new natural product-inspired optimization campaigns before outcomes are known. Such studies would need transparent candidate selection, predefined decision criteria, documentation of held and rejected candidates, and reporting of negative results. Without this structure, post hoc selection could make the model appear more useful than it actually is. The model should therefore be evaluated not only by whether it identifies attractive candidates, but also by whether it prevents unjustified progression.

Several failure modes deserve explicit attention in future assessment. Random-split inflation, repeated scaffold families, unavailable negative data, assay-context mismatch, synthetic-route failure, uncertainty miscalibration, and confirmation bias can all make a computational workflow appear more reliable than it is. The proposed model is intended to expose these risks rather than eliminate them automatically. Its validation should therefore test whether the architecture improves transparency, evidence sufficiency, and decision discipline before any claim of practical superiority is made.

CONCLUSION

Lead optimization in natural product chemical space cannot be reduced to generating valid molecules or maximizing predicted activity. Natural product-inspired candidates often combine structural richness with prediction uncertainty, synthetic difficulty, and pharmacological complexity. A credible optimization workflow must therefore treat chemical-space representation, generative design, activity prediction, feasibility analysis, plausibility assessment, and uncertainty handling as interacting evidence layers.

The methodological model proposed here provides a structured way to keep those evidence layers connected without hiding conflict inside a single unexplained score. It distinguishes novelty from value, predicted activity from pharmacological plausibility, synthetic accessibility from practical route feasibility, and computational ranking from experimentally credible progression. Its central decision logic is conditional: candidates may be prioritized, optimized, redesigned, held, or rejected depending on the complete evidence pattern rather than on any isolated computational output.

The article's contribution is a methodological architecture for natural product lead optimization that emphasizes constrained generation, evidence sufficiency, cross-domain conflict resolution, and selective redesign. The model does not claim prospective validation, improved discovery success, or superiority over existing workflows. Its value lies in clarifying what should be demonstrated before a natural product-derived or natural product-inspired candidate is treated as ready for experimental progression. Future work should test the model empirically through transparent retrospective and prospective applications.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

REFERENCES

[1] Pye CR, Bertin MJ, Lokey RS, Gerwick WH, Linington RG. Retrospective analysis of natural products provides insights for future discovery trends. Proc Natl Acad Sci U S A. 2017;114(22):5601-6. doi:10.1073/pnas.1614680114

[2] Atanasov AG, Zotchev SB, Dirsch VM, International Natural Product Sciences Taskforce, Supuran CT. Natural products in drug discovery: Advances and opportunities. Nat Rev Drug Discov. 2021;20(3):200-16. doi:10.1038/s41573-020-00114-z

[3] Stone S, Newman DJ, Colletti SL, Tan DS. Cheminformatic analysis of natural product-based drugs and chemical probes. Nat Prod Rep. 2022;39(1):20-32. doi:10.1039/D1NP00039J

[4] Chen Y, Kirchmair J. Cheminformatics in natural product-based drug discovery. Mol Inform. 2020;39(12). doi:10.1002/minf.202000171

[5] Grisoni F, Merk D, Consonni V, Hiss JA, Giani Tagliabue S, Todeschini R, et al. Scaffold hopping from natural products to synthetic mimetics by holistic molecular similarity. Commun Chem. 2018;1:44. doi:10.1038/s42004-018-0043-x

[6] Olivecrona M, Blaschke T, Engkvist O, Chen H. Molecular de-novo design through deep reinforcement learning. J Cheminform. 2017;9(1):48. doi:10.1186/s13321-017-0235-x

[7] 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. doi:10.1021/acscentsci.7b00572

[8] Zhavoronkov A, Ivanenkov YA, Aliper A, Veselov MS, Aladinskiy VA, Aladinskaya AV, et al. Deep learning enables rapid identification of potent DDR1 kinase inhibitors. Nat Biotechnol. 2019;37(9):1038-40. doi:10.1038/s41587-019-0224-x

[9] Krenn M, Häse F, Nigam AK, Friederich P, Aspuru-Guzik A. Self-referencing embedded strings (SELFIES): A 100% robust molecular string representation. Mach Learn Sci Technol. 2020;1(4):045024. doi:10.1088/2632-2153/aba947

[10] Meyers J, Fabian B, Brown N. De novo molecular design and generative models. Drug Discov Today. 2021;26(11):2707-15. doi:10.1016/j.drudis.2021.05.019

[11] Korshunova M, Huang N, Capuzzi S, Radchenko DS, Savych O, Moroz YS, et al. Generative and reinforcement learning approaches for the automated de novo design of bioactive compounds. Commun Chem. 2022;5(1):129. doi:10.1038/s42004-022-00733-0

[12] Wu Z, Ramsundar B, Feinberg EN, Gomes J, Geniesse C, Pappu AS, et al. MoleculeNet: A benchmark for molecular machine learning. Chem Sci. 2018;9(2):513-30. doi:10.1039/C7SC02664A

[13] Lenselink EB, ten Dijke N, Bongers B, Papadatos G, van Vlijmen HWT, Kowalczyk W, et al. Beyond the hype: Deep neural networks outperform established methods using a ChEMBL bioactivity benchmark set. J Cheminform. 2017;9(1):45. doi:10.1186/s13321-017-0232-0

[14] 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. doi:10.1021/acs.jcim.9b00237

[15] Feinberg EN, Sur D, Wu Z, Husic BE, Mai H, Li Y, et al. PotentialNet for molecular property prediction. ACS Cent Sci. 2018;4(11):1520-30. doi:10.1021/acscentsci.8b00507

[16] Torng W, Altman RB. Graph convolutional neural networks for predicting drug-target interactions. J Chem Inf Model. 2019;59(10):4131-49. doi:10.1021/acs.jcim.9b00628

[17] Soleimany AP, Amini A, Goldman S, Rus D, Bhatia SN, Coley CW. Evidential deep learning for guided molecular property prediction and discovery. ACS Cent Sci. 2021;7(8):1356-67. doi:10.1021/acscentsci.1c00546

[18] Coley CW, Rogers L, Green WH, Jensen KF. Computer-assisted retrosynthesis based on molecular similarity. ACS Cent Sci. 2017;3(12):1237-45. doi:10.1021/acscentsci.7b00355

[19] Segler MHS, Preuss M, Waller MP. Planning chemical syntheses with deep neural networks and symbolic AI. Nature. 2018;555(7698):604-10. doi:10.1038/nature25978

[20] Schwaller P, Laino T, Gaudin T, Bolgar P, Hunter CA, Bekas C, et al. Molecular transformer: A model for uncertainty-calibrated chemical reaction prediction. ACS Cent Sci. 2019;5(9):1572-83. doi:10.1021/acscentsci.9b00576

[21] Gao W, Coley CW. The synthesizability of molecules proposed by generative models. J Chem Inf Model. 2020;60(12):5714-23. doi:10.1021/acs.jcim.0c00174

[22] 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. doi:10.1039/D0SC05401A

[23] 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). doi:10.1093/nar/gkab255

[24] Bender A, Cortés-Ciriano I. Artificial intelligence in drug discovery: What is realistic, what are illusions? Part 1: Ways to make an impact, and why we are not there yet. Drug Discov Today. 2021;26(2):511-24. doi:10.1016/j.drudis.2020.12.009

[25] Bender A, Cortés-Ciriano I. Artificial intelligence in drug discovery: What is realistic, what are illusions? Part 2: A discussion of chemical and biological data. Drug Discov Today. 2021;26(4):1040-52. doi:10.1016/j.drudis.2020.11.037

[26] Chaudhari R, Tan Z, Huang B, Zhang S. Computational polypharmacology: A new paradigm for drug discovery. Expert Opin Drug Discov. 2017;12(3):279-91. doi:10.1080/17460441.2017.1280024

[27] Pinzi L, Rastelli G. Molecular docking: Shifting paradigms in drug discovery. Int J Mol Sci. 2019;20(18):4331. doi:10.3390/ijms20184331

[28] Wallach I, Heifets A. Most ligand-based classification benchmarks reward memorization rather than generalization. J Chem Inf Model. 2018;58(5):916-32. doi:10.1021/acs.jcim.7b00403

[29] Bosc N, Atkinson F, Felix E, Gaulton A, Hersey A, Leach AR. Large scale comparison of QSAR and conformal prediction methods and their applications in drug discovery. J Cheminform. 2019;11(1):4. doi:10.1186/s13321-018-0325-4