Pharmacophore

ISSN-2229-5402

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

SYNTHESIS BEGINS AT GENERATION: EMBEDDING ROUTE FEASIBILITY, MATERIAL AVAILABILITY, AND PROCESS BURDEN INTO MOLECULAR DESIGN

Michael Roberts[1]*, Sarah Thompson[2], James Anderson[3], Robert Smith[4]

1. ***Department of Synthesis-Aware Molecular Generation, Faculty of Pharmacy, University of Glasgow, Glasgow, United Kingdom.***
2. ***Department of Route Feasibility and Material Availability, Faculty of Pharmacy, University of Auckland, Auckland, New Zealand.***
3. ***Department of Process Burden and Manufacturability, Faculty of Pharmacy, University of Manchester, Manchester, United Kingdom.***
4. ***Department of Synthesis-Embedded Molecular Design, Faculty of Pharmacy, University of Edinburgh, Edinburgh, United Kingdom.***

ARTICLE INFO

Received:
17 July 2025
Received in revised form:
27 November 2025
Accepted:
27 November 2025
Available online:
28 December 2025

Keywords: Generative molecular design, Retrosynthesis, Synthetic accessibility, Route feasibility, Material availability, Process burden

ABSTRACT

Generative molecular design can propose chemically valid structures with attractive predicted properties, yet many candidates remain disconnected from a credible path to synthesis. Treating synthesizability as a post hoc filter creates a structural inefficiency: molecular objectives are optimized before route feasibility, precursor availability, operational burden, and planning uncertainty are allowed to influence the design itself. This article develops an original synthesis-constrained design theory in which synthesis begins during molecular generation rather than after it. The proposed construct, termed Route-Embedded Design Viability, represents each candidate as a coupled molecule–route hypothesis characterized by molecular value, step-level reaction feasibility, route topology, context-specific material availability, anticipated process burden, and an explicit uncertainty envelope. The theory distinguishes molecular plausibility from route viability, predicted reaction feasibility from experimental confirmation, catalogue presence from usable material access, and early process indicators from validated manufacturing evidence. It further proposes that molecular value and route viability should be optimized jointly through constrained or Pareto-based search rather than collapsed automatically into a single score. Evaluation should therefore progress from structural integrity and predictive relevance to full retrosynthetic search, route comparison, material verification, process-burden assessment, uncertainty analysis, and expert review of an executable synthesis plan. The contribution is conceptual and methodological rather than empirically validated. Its applicability remains conditional on reaction-data quality, planner configuration, stock definitions, process assumptions, project objectives, and human expertise. By repositioning synthesis evidence as a generative input, the theory provides a bounded basis for designing molecules whose predicted pharmaceutical value is considered together with the practical consequences of making them.

To Cite This Article: Roberts M, Thompson S, Anderson J, Smith R. Synthesis Begins at Generation: Embedding Route Feasibility, Material Availability, and Process Burden into Molecular Design. Pharmacophore. 2025;16(6):56-65. https://doi.org/10.51847/5QN65SjElH

Introduction

Generative molecular design reframes discovery as an inverse problem in which computational systems search chemical space for structures expected to satisfy specified objectives. This formulation has enabled diverse architectures for proposing molecules, optimizing predicted properties, and navigating large combinatorial design spaces. However, pharmaceutical relevance cannot be inferred from molecular generation alone. A candidate becomes useful only when its predicted value can be interpreted within medicinal chemistry, experimental, and project-decision contexts, while the choice of representation, generative architecture, and optimization objective determines which structures are reachable and preferentially explored [1–

Corresponding Author: Michael Roberts; Department of Synthesis-Aware Molecular Generation, Faculty of Pharmacy, University of Glasgow, Glasgow, United Kingdom. E-mail: michael.roberts@glasgow.ac.uk

3]. The central difficulty is therefore not merely whether a model can generate a valid molecule, but whether the design process can preserve pharmaceutical value while remaining connected to credible chemical realization.

Representations and objective functions encode assumptions about what constitutes a desirable candidate. Graphs, strings, latent vectors, reaction actions, and fragment-based representations expose different structural relationships and support different forms of search. Likewise, property predictors, similarity functions, novelty terms, and synthetic-accessibility estimates define different optimization landscapes. These choices influence the distribution of generated structures and can silently exclude or favour particular chemical regions [4]. When synthesis is represented only through a late scalar penalty, the generator receives little information about why a candidate is difficult to make, whether a different route might resolve the difficulty, or whether a small structural revision could retain molecular value while improving route viability.

The gap becomes visible when molecular plausibility is compared with retrosynthetic accessibility. A structure may satisfy formal valence rules, resemble known drug-like molecules, and receive favourable predicted-property scores while remaining unsolved by a synthesis planner or requiring routes that are implausible under the relevant resource constraints. Comparative analysis of generated molecules has shown that structural validity and common synthetic-accessibility measures do not guarantee that a retrosynthetic system can identify a credible route [5]. Even a planner-generated route is only a computational hypothesis: its value depends on the reaction model, search policy, available precursor stock, route-scoring method, and the extent to which predicted transformations correspond to experimentally usable chemistry.

This article proposes that synthesis should be treated as a constitutive part of molecular generation rather than as a downstream accept–reject operation. The proposed theory introduces Route-Embedded Design Viability, or REDV, as an uncertainty-qualified relationship between a molecular proposal and at least one contextually plausible route hypothesis. REDV does not define synthesizability as a fixed intrinsic property of a molecule. Instead, it organizes molecular value, reaction feasibility, route complexity, material availability, process burden, and uncertainty as interacting dimensions that must be examined together. The objective is not to present a validated model or universal scoring system, but to establish a conceptual architecture for designing and evaluating molecule–route pairs without conflating computational planning with experimental synthesis, process development, manufacturability, or therapeutic utility.

Why Post Hoc Synthesis Filtering Is Structurally Inefficient

Widely used molecular-generation benchmarks assess validity, uniqueness, novelty, diversity, distributional similarity, and performance against goal-directed objectives. These measures are useful for identifying malformed outputs, mode collapse, lack of novelty, and failure to optimize declared targets. They do not, however, establish whether a proposed structure can be connected to an executable route, obtained from accessible materials, or prepared without disproportionate operational burden. Moreover, favourable benchmark performance can coexist with chemically unhelpful behaviour when a generator exploits scoring functions, concentrates on narrow structural motifs, or satisfies a numerical objective through solutions that do not correspond to meaningful discovery opportunities [6–8]. Post hoc synthesis filtering inherits this benchmark logic by allowing the generator to optimize a molecule-centred objective first and asking whether the result can be made only after computational effort has already been concentrated on it.

The inefficiency is structural because the information needed to avoid many failures is, in principle, available during generation. A retrosynthesis model may indicate that a structural motif lacks plausible disconnections, that the route depends on poorly supported transformations, or that all identified routes terminate in inaccessible precursors. A material database may reveal that the necessary building blocks are unavailable in the relevant laboratory or region. A preliminary process assessment may identify excessive protecting-group operations, hazardous reagents, difficult separations, or a route architecture likely to impose high material intensity. When these signals are withheld until the end, the generator cannot revise the molecular structure in response. It can only produce a candidate that is subsequently accepted or discarded. This separation also increases vulnerability to objective exploitation, because a generator may learn to maximize predicted molecular properties without receiving a corresponding penalty for route fragility or process burden [9].

Post hoc filtering also compresses heterogeneous evidence into a decision that appears simpler than the underlying problem. A candidate may be rejected because no route is found, but this result can have several explanations: the molecule may genuinely be difficult to synthesize; the reaction model may lack relevant precedents; the search procedure may fail to explore an appropriate route family; the allowed stock may be too restrictive; or the planner may rank routes using assumptions misaligned with the project. Conversely, a candidate may pass a synthetic-accessibility threshold while depending on an impractical sequence, unavailable catalyst, unstable intermediate, or burdensome purification. Evaluation frameworks that combine multiple molecular scoring functions and chemistry diagnostics demonstrate that conclusions depend on the selected objectives, task definitions, and analytical checks [10]. The same principle applies more strongly to synthesis, where the relevant evidence is distributed across molecules, reactions, routes, materials, operating assumptions, and organizational context.

REDV addresses this structural separation by replacing the sequence “generate, score, filter” with a coupled relation between molecular and route hypotheses. During generation, a candidate is provisionally associated with one or more route representations, and evidence from reaction feasibility, route topology, material availability, and process burden is returned to the design process. Failure does not automatically imply rejection; it can trigger a specific revision. Unsupported transformations may prompt route exploration, unavailable materials may prompt precursor substitution, excessive burden may prompt scaffold modification, and high uncertainty may prompt evidence acquisition rather than optimization. This

proposed feedback structure extends the observation that generated molecular validity does not ensure retrosynthetic accessibility [5]. It does not claim that every route liability can be anticipated or that synthesis-aware generation will necessarily outperform sequential workflows. Its purpose is to ensure that synthesis-relevant evidence can influence the molecule before the design becomes committed to a structurally attractive but operationally disconnected solution. **Table 1** organizes the evidence, constructs, and interpretation boundaries needed to develop why post hoc synthesis filtering is structurally inefficient within the article's central argument.

Table 1. Evidence domains, core questions, scientific requirements, and interpretation boundaries for Why post hoc synthesis filtering is structurally inefficient

Construct or evidence domain	Core question	Scientific basis	Role in the article	Failure or overclaiming risk	Boundary statement
Molecular validity	Is the generated structure chemically interpretable under the chosen representation?	Valence, connectivity, representation integrity, and structural-quality checks	Establishes the minimum condition for evaluating a generated proposal	Validity may be misrepresented as evidence of synthesis feasibility	A valid molecular representation is not necessarily synthesizable, developable, safe, or therapeutically useful
Predicted molecular value	Does the candidate satisfy the project's declared molecular objectives?	Activity, selectivity, physicochemical, exposure, safety, or other task-specific predictions	Explains why the candidate merits continued consideration	Optimization may exploit imperfect predictors or leave their applicability domain	Predicted value is not equivalent to experimental activity, mechanism, or clinical utility
Post hoc synthetic-accessibility filter	Does a late-stage score or planner classify the completed molecule as accessible?	Structural heuristics, learned scores, reaction-derived proxies, or retrosynthesis outputs	Represents the conventional sequential design pattern criticized in the theory	A pass–fail decision can hide why the molecule or route is difficult	A late filter is a triage device, not a complete representation of route viability
Step-level reaction feasibility	Are proposed transformations supported by relevant precedents, substrates, and conditions?	Reaction corpora, templates, learned models, precedent retrieval, and expert assessment	Identifies where a route hypothesis is chemically weak	Model confidence may be treated as experimental success probability	Predicted feasibility remains conditional until reactions are experimentally tested
Route topology and complexity	Is the complete route strategically manageable rather than merely discoverable?	Step count, longest linear sequence, convergence, branching, protecting groups, and dependencies	Distinguishes nominal route completion from route quality	Simple metrics may reward short but fragile or unrealistic routes	No route descriptor alone establishes practical synthesis
Material availability	Are required starting materials, reagents, catalysts, and enabling resources accessible in the declared context?	Internal inventory, supplier records, specifications, quantities, and lead times	Makes route termination and resource assumptions explicit	Catalogue listing may be mistaken for actual availability or suitability	Availability is time-, region-, organization-, quantity-, and specification-dependent
Process burden	What material, operational, safety, purification, and scale burdens accompany the route?	Preliminary route features, process-mass indicators, hazard information, and unit-operation demands	Extends synthesis assessment beyond reaction plausibility	Early estimates may be presented as validated process economics or manufacturability	Process-burden indicators are provisional and cannot replace process development
Uncertainty and provenance	How reliable, calibrated, and contextually applicable is each synthesis-related claim?	Calibration, model disagreement, data coverage, route alternatives, and evidence lineage	Prevents deterministic interpretation of computational planning	Confidence scores may conceal domain shift or correlated errors	Uncertainty must be represented at molecule, reaction, route, material, and process levels
Coupled molecule–route revision	Can synthesis evidence modify the molecule, route, or both before final selection?	Iterative generation, route feedback, constrained search, and expert review	Defines the structural alternative to post hoc filtering	The generator may overfit to narrow reaction spaces or favour familiar chemistry	Coupling supports revision but does not guarantee superior molecular or experimental outcomes

Reaction Feasibility, Route Complexity, Materials, and Process Burden

Reaction feasibility and route viability are related but non-equivalent. Reaction-informed complexity models can estimate whether a product appears more synthetically advanced than its likely precursors, while multistep planners recursively apply learned or encoded disconnections to connect a target with candidate starting materials. Open planning systems further demonstrate that route outcomes depend on reaction policies, tree-search procedures, route scores, and the definition of terminal stock [11–13]. These approaches provide essential evidence, but they address different units of analysis. A complexity score describes a molecular or transformation-related tendency; a single-step model proposes disconnections; a multistep planner constructs route trees; and a stock determines when the search is considered complete. REDV therefore treats reaction feasibility, route topology, and material termination as separate components rather than interchangeable measures of synthesizability.

Step-level reaction feasibility concerns whether a proposed transformation is chemically plausible for the specific substrates, products, and conditions under consideration. Evidence may include reaction precedents, template applicability, model calibration, condition prediction, mechanistic compatibility, and expert interpretation. Route complexity concerns how feasible steps are arranged into a complete strategy. A route can consist of individually plausible transformations yet remain unattractive because it is excessively linear, dependent on unstable intermediates, protecting-group intensive, poorly convergent, or vulnerable to failure at a late stage. Process-aware route design further shows that synthetic choices influence material use and environmental burden and that these consequences should be considered before a route becomes entrenched [14]. Accordingly, REDV represents route viability as a structured profile rather than as the product of step probabilities or a shortest-path calculation.

Material availability adds a contextual terminal condition to route planning. A route that reaches a commercially catalogued precursor may still be unusable if the material is unavailable in the required region, quantity, purity, stereochemical form, or delivery period. Internal inventories impose a different boundary from broad commercial catalogues, and laboratory capabilities determine whether specialist transformations, purification methods, or containment requirements are realistic. Material availability should therefore be represented with time-stamped provenance and at least basic specification information, rather than as a permanent binary label. Process burden is similarly contextual. Structure-based prediction may support an early estimate of process mass intensity, but such an estimate cannot substitute for route-specific evidence on solvents, reagents, reaction concentration, work-up, isolation, purification, safety, waste, and scale behaviour [15]. Its legitimate role is to flag likely burden early enough to influence design, not to certify a manufacturing route.

The proposed REDV construct integrates these dimensions without assuming that they can be reduced to one universal score. Reaction feasibility may be expressed at the step level; route complexity through graph and strategy descriptors; material availability through verified stock states; and process burden through a provisional vector of material, operational, safety, and separation demands. These dimensions may then be subject to hard constraints, project-specific thresholds, or Pareto comparison. A route dependent on one uncertain transformation should not necessarily be ranked identically to a route with many moderately supported steps, even when their aggregate scores are similar. Likewise, a short route based on unavailable materials may be less viable than a longer route using established chemistry and accessible precursors. The theory therefore treats synthesis viability as conditional on a declared route–resource context and requires that every favourable conclusion retain its assumptions, uncertainty, and evidentiary limits.

Embedding Synthetic Constraints into Generative Representations

Synthetic constraints can influence molecular design only when the representation exposes actions or features that connect structural modification to chemical realization. One approach is to define the generative action space through reaction rules, fragments, or transformations rather than unrestricted atom-level editing. Hybrid systems combining machine intelligence with rule-driven synthesis demonstrate that generated structures can be conditioned on permitted chemical operations [16]. This restriction changes the meaning of generation: the system no longer proposes an isolated molecular graph and subsequently asks how it might be made, but constructs the proposal through a sequence of transformations that can also serve as an initial route hypothesis. The resulting feasibility nevertheless remains conditional on the scope and correctness of the encoded reaction rules.

Experimentally bounded platforms provide a stronger but narrower form of synthesis embedding. Generative design has been coupled with on-chip synthesis in a defined reaction space, showing how a model can search among structures designed around available experimental operations [17]. Such coupling reduces the distance between a generated candidate and its physical realization because the representation reflects the platform's actual chemistry. It also illustrates an important boundary: feasibility within one synthesis platform cannot be generalized to chemistry beyond its supported transformations, substrates, equipment, or operating conditions. Platform-conditioned generation is therefore an example of context-specific REDV rather than evidence for universally executable molecular design.

Reaction-based scaffold decoration offers another architecture in which building blocks and transformations define permissible molecular changes. LibINVENT, for example, generates decorated scaffolds through reaction-aware actions intended to preserve medicinal-chemistry relevance and synthetic grounding [18]. Within REDV, this architecture can be understood as a constrained design grammar. The grammar links each structural modification to a prospective chemical action and an associated material requirement. Its value depends on reaction-template coverage, compatibility between templates and substrates, the quality of the building-block collection, and whether the resulting transformation remains plausible under the

intended laboratory conditions. Restricting the grammar may improve route interpretability while simultaneously narrowing chemical novelty.

Synthetic constraints may also be introduced as evaluative feedback rather than as hard restrictions. Retrosynthesis-derived accessibility signals can be incorporated during optimization so that route evidence affects which molecules are rewarded and retained [19]. More directly, a retrosynthesis planner can function as an optimization oracle, allowing sample-efficient reinforcement learning to favour molecules for which route hypotheses can be identified [20]. These approaches move beyond post hoc filtering because synthesis information changes the generative trajectory. However, they also transmit planner assumptions into the molecular distribution. A model may learn the biases of the reaction corpus, favour familiar transformations, exploit weaknesses in the route score, or generate structures tailored to an incomplete stock definition. REDV therefore requires synthesis feedback to retain its provenance and uncertainty rather than becoming an unqualified reward.

Joint Optimization of Molecular Value and Route Viability

Molecular value and route viability are naturally multi-dimensional. A candidate may improve predicted potency while worsening selectivity, solubility, safety, route complexity, material access, or process burden. Continuous latent-space optimization has shown how multiple molecular properties can be pursued simultaneously, but such methods generally treat synthesis as one optional objective rather than a structured route-dependent state [21]. REDV extends multi-objective design by defining the optimization unit as a molecule–route pair. Each candidate may be associated with several routes, and each route may produce a different profile of reaction support, precursor access, operational burden, and uncertainty. The same molecule can therefore occupy different viability positions under different route and resource contexts.

Pareto-based optimization offers a suitable logic for preserving these trade-offs. Rather than converting every objective into one weighted sum, non-dominated ranking retains candidates for which no alternative is superior across all declared dimensions [22]. In synthesis-constrained design, a Pareto set might include a highly promising molecule with a demanding but credible route, a moderately promising molecule with a robust and accessible route, and a chemically novel candidate whose route uncertainty justifies further investigation. This representation makes trade-offs visible to chemists and project teams. It also avoids the misleading implication that a universal coefficient can determine how much predicted biological value should compensate for unavailable materials or severe process liabilities.

Computational cost complicates joint optimization because full retrosynthetic searches, route comparisons, and process evaluations are more expensive than many molecular-property calculations. Sample-efficient reinforcement-learning strategies can reduce the number of oracle evaluations needed to guide molecular optimization [23]. REDV would use staged evidence acquisition: inexpensive molecular checks and rapid synthesis proxies may screen early proposals, while detailed route planning and process assessment are reserved for candidates that remain competitive. This hierarchy should not be interpreted as permission to substitute a proxy permanently for a full evaluation. Its purpose is to allocate computational and expert effort while preserving a clear distinction between preliminary triage and route-level evidence.

Synthesis-planning-driven optimization can further connect property-directed molecular modification with reaction templates and reference routes [24]. Within the proposed theory, joint optimization becomes an iterative negotiation among molecular objectives, route alternatives, resource constraints, and uncertainty. A molecular modification may be accepted because it improves value without materially weakening the route; rejected because it introduces an unsupported transformation; or retained provisionally because a second route reduces the initial liability. Hard constraints may be appropriate for prohibited materials or safety-critical operations, whereas Pareto comparison may be preferable for route length, cost, environmental burden, or molecular performance. REDV does not prescribe universal thresholds. It requires that thresholds be explicit, project-specific, and reviewable.

Uncertainty in Retrosynthesis and Manufacturing Assumptions

Retrosynthetic predictions are uncertain at several levels. A single-step model may assign high probability to a disconnection even though the specific substrates, stereochemical demands, or reaction conditions fall outside the supporting evidence. The Molecular Transformer demonstrated that reaction prediction can incorporate uncertainty calibration, emphasizing that prediction confidence should be evaluated rather than assumed [25]. For REDV, step-level uncertainty is not a decorative confidence score. It determines whether a transformation should support route selection, trigger alternative-route exploration, or be treated as an unresolved experimental question. Uncalibrated model scores must not be interpreted as probabilities of laboratory success.

Dataset construction creates another source of uncertainty. Computer-assisted synthesis-planning systems depend on reaction corpora whose coverage, duplication, reporting quality, template specificity, and historical biases influence the transformations that models can learn [26]. Patent-derived data may overrepresent successful, disclosed, or commercially significant chemistry while underrepresenting failures, routine operations, and tacit procedural knowledge. Internal pharmaceutical datasets may provide different coverage but remain organization-specific. REDV therefore requires provenance at the reaction-policy level, including the data domain, template generation method, stock definition, and principal exclusions. A route should be interpreted as conditional on this evidence environment rather than as an intrinsic property of the target molecule.

Uncertainty also propagates from single-step predictions into multistep search. Comparative work has shown that changing the one-step retrosynthesis model can alter route-solving outcomes and the routes recovered by a planner [27]. Search heuristics may amplify early prediction errors by allocating computational effort to one route family while neglecting alternatives. Route-

level confidence therefore cannot be estimated reliably by examining only the highest-ranked transformation at each step. REDV instead calls for sensitivity analysis across models, search settings, stock definitions, and route alternatives. Agreement across configurations can strengthen confidence, although agreement among systems trained on similar data may still reflect shared bias rather than independent confirmation.

Manufacturing assumptions add uncertainty that retrosynthesis models rarely represent fully. A route may appear plausible while lacking evidence about reagent quality, heat and mass transfer, mixing, impurity formation, isolation, waste handling, equipment compatibility, or scale-dependent safety. Reliability and interpretability analyses of machine-learning retrosynthesis reinforce the need to examine model behaviour beyond nominal ranking performance [28]. REDV extends that caution to process claims. Early process-burden estimates should be expressed as provisional flags tied to stated assumptions. They can support route comparison and identify questions for process chemists, but they cannot establish scalability, cost, environmental performance, or manufacturability without route-specific experimental development.

Evaluation from Generated Structure to Executable Synthesis Plan

Evaluation should begin with the generated structure but must not end there. Rapid planner-derived classifiers can approximate whether a full retrosynthetic search is likely to solve a molecule, making them useful for high-throughput triage [29]. Their predictions remain conditional on the planner, stock, reaction policies, and labels from which they were learned. REDV therefore places rapid synthesizability proxies below full route analysis in an evidence hierarchy. A favourable proxy result permits deeper evaluation; it does not establish that a route exists, that the necessary materials are available, or that the chemistry can be executed under relevant conditions.

Full route evaluation requires a representation capable of comparing synthetic plans at the route level. LinChemIn provides a common graph-based model for converting, analysing, and comparing routes generated by different planning systems [30]. Such representations allow route topology, reaction sequences, branching, convergence, and shared intermediates to be examined systematically. Within REDV, route graphs also provide the structure needed to attach step-level uncertainty, material states, process flags, and provenance. Route comparison should remain plural whenever possible, because a single highest-ranked plan may conceal strategic alternatives that are more robust, accessible, or compatible with the project's operational context.

Practical experience with synthesis-planning systems indicates that planner assessment should include more than the proportion of targets classified as solved. Route shape, stock utilization, reaction-space coverage, target-set composition, and unintended exploitation of templates can reveal behaviours hidden by aggregate success measures [31]. Evaluation should therefore test whether synthesis-embedded generation merely produces molecules tailored to planner weaknesses or genuinely improves the quality of molecule–route alternatives. Relevant comparisons include post hoc filtering, hard reaction constraints, rapid synthesis proxies, full planner feedback, and joint molecule–route optimization. Computational superiority alone would remain insufficient; prospective utility requires evidence that the proposed outputs improve expert decisions or experimental outcomes.

The highest pre-experimental REDV stage is an expert-reviewable executable synthesis plan. Such a plan requires a complete route hypothesis, verified material states, transformation evidence, relevant conditions, process and safety flags, alternative strategies, and an explicit uncertainty record. Work on generation constrained by in-house resources demonstrates that changing the permitted inventory can materially alter which molecules and routes are considered accessible [32]. This confirms that executable-plan status is contextual rather than molecularly intrinsic. **Table 2** organizes the evidence, constructs, and interpretation boundaries needed to develop evaluation from generated structure to executable synthesis plan within the article's central argument.

Table 2. Evaluation, implementation, and research priorities arising from Synthesis Begins at Generation Embedding Route Feasibility, Material Availability, and Process Burden

Component or layer	Inputs	Core function	Expected output	Uncertainty or limitation	Validation requirement
Structural-integrity layer	Molecular representation, valence checks, duplicate analysis, domain descriptors	Confirm that the generated proposal is chemically interpretable and relevant to the intended design space	Valid molecular candidate for further evaluation	Representation validity does not imply stability or synthesizability	Expert and computational review across diverse molecular classes
Molecular-value layer	Task-specific property models, applicability evidence, uncertainty estimates	Determine why the molecule merits consideration	Multi-dimensional predicted-value profile	Predictor error, domain shift, and oracle exploitation	Prospective experimental testing of prioritized and rejected candidates
Rapid synthesis-triage layer	Synthetic-complexity or planner-derived accessibility proxies	Allocate detailed planning effort efficiently	Preliminary synthesis-priority category	Proxy inherits the assumptions and errors of its training labels	Comparison with full planning and expert assessment

Reaction-feasibility layer	Reaction models, templates, precedents, substrates, conditions, calibration	Assess the plausibility of individual transformations	Step-level feasibility and uncertainty record	Recorded precedent may not transfer to the proposed context	Prospective execution of representative high- and low-confidence reactions
Multistep route layer	Search algorithm, reaction policy, stock definition, route score	Connect the target to declared starting materials	One or more complete route hypotheses	Search failure, policy bias, incomplete route exploration	Cross-planner and cross-configuration replication
Route-comparison layer	Common route graphs, topology measures, strategic descriptors	Compare route families rather than isolated scores	Ranked or Pareto-organized route alternatives	Metrics may omit tacit chemistry and operational difficulty	Blinded comparison with synthetic-chemist judgments
Material-availability layer	Internal inventory, suppliers, quantity, purity, stereochemical form, lead time	Verify that route inputs are accessible in the stated context	Time-stamped resource-availability profile	Supplier and inventory states change over time	Live inventory checks and multi-site transportability studies
Process-burden layer	Reagents, solvents, unit operations, hazards, purification, scale assumptions	Identify early operational and material liabilities	Provisional process-risk vector	Sparse early data and unmodelled scale effects	Retrospective and prospective comparison with developed processes
Uncertainty and sensitivity layer	Calibration, model disagreement, alternative routes, stock and policy changes	Determine the stability of viability conclusions	Assumption-qualified confidence profile	Correlated errors and shared dataset biases	Empirical coverage tests and route-level uncertainty studies
Human-review layer	Complete evidence record, project objectives, chemistry expertise, safety and resource constraints	Decide whether experimental escalation is justified	Expert-reviewed executable synthesis plan	Judgment varies by expertise, organization, and project priorities	Structured multi-expert review and prospective decision-utility studies
Experimental-confirmation layer	Executed reactions, analytical evidence, failures, process observations	Test whether the proposed route functions under defined conditions	Context-specific synthesis evidence	Success at one scale or site may not transfer	Repetition, route refinement, and appropriately staged process development

Figure 1 presents the synthesis-embedded molecular generation loom, showing how the article's key components and boundaries are connected within evaluation from generated structure to executable synthesis plan.

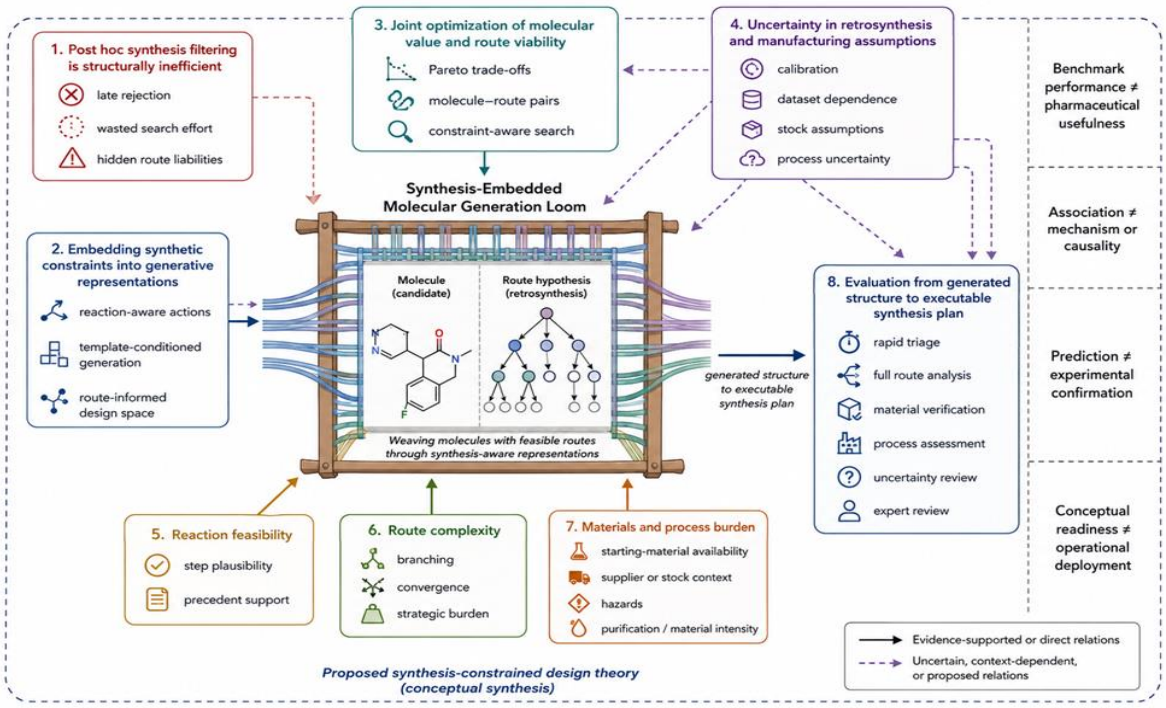

Figure 1. Synthesis-Embedded Molecular Generation Loom.

The figure is an original conceptual synthesis that organizes the article's central contribution across post hoc synthesis filtering is structurally inefficient, reaction feasibility, route complexity, materials, and process burden, reaction feasibility, route complexity, process burden, embedding synthetic constraints into generative representations. Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, regulatory determination, clinical recommendation, or deployment-ready system.

Limitations and Boundary Conditions

REDV is a proposed theory rather than a validated predictive architecture. Its components draw on evidence from generative design, reaction prediction, retrosynthesis, route analysis, material inventories, and early process assessment, but these domains differ in data maturity, validation practice, and unit of analysis. Combining them conceptually does not establish that their uncertainties can be aggregated reliably or that joint optimization will improve experimental decisions. Reaction models may be calibrated within recorded datasets while remaining poorly applicable to novel substrates, and process proxies may identify broad associations without capturing route-specific operating behaviour. The framework therefore organizes evidence requirements but does not supply missing evidence.

The construct is also strongly context-dependent. A molecule may be route-viable under one reaction policy, stock definition, supplier environment, laboratory capability, or process-risk tolerance and non-viable under another. Material inventories change, commercial listings become obsolete, and specialist chemistry may be accessible to one organization but unavailable to another. Similarly, project priorities determine whether molecular value justifies route complexity or whether a resource constraint should function as a hard exclusion. REDV should not be used to conceal these decisions inside fixed weights. Every viability claim must identify the route–resource context, evidence provenance, uncertainty, and human decision criteria on which it depends.

Misuse could arise if synthesis embedding is interpreted as permission to automate experimental escalation or to equate planner success with manufacturability. The proposed architecture cannot establish reaction yield, impurity control, scale-up safety, economic feasibility, environmental superiority, developability, clinical utility, or regulatory acceptability. It may also bias generation toward historically common chemistry, well-represented reaction classes, or readily purchased building blocks, thereby suppressing innovative routes and underexplored molecular regions. Human review is consequently not a temporary substitute for future automation but a constitutive safeguard for interpreting incomplete evidence, recognizing exceptional chemistry, and deciding when uncertainty warrants experimentation rather than rejection.

Research Agenda and Implementation Priorities

The first research priority is to evaluate whether synthesis-embedded generation improves decisions relative to sequential generation and filtering. Controlled comparisons should use matched molecular objectives and distinguish hard reaction constraints, rapid synthesis proxies, full retrosynthesis feedback, and joint molecule–route optimization. Evaluation should examine molecular diversity, route robustness, material accessibility, process flags, uncertainty calibration, and expert decisions rather than relying only on solved-route rates. Prospective studies should document both successful and failed transformations and should test whether candidates prioritized by REDV produce more useful experimental choices, not merely more favourable computational scores.

A second priority is the development of interoperable evidence infrastructure. Route graphs should carry reaction provenance, condition information, model version, stock definition, material specification, process flags, and uncertainty at each level. Time-stamped inventories and supplier records are needed to distinguish theoretical commercial availability from usable access. Reaction datasets should better represent negative results, condition dependence, and chemistry outside heavily patented domains. Process-development records could support interpretable early burden models, provided that confidentiality, dataset shift, and scale dependence are addressed. Shared reporting standards would allow route hypotheses from different planning systems to be compared without implying that all route metrics are equally mature.

Implementation should proceed in stages. Initial use should be retrospective, examining how REDV components would have altered decisions in completed projects. The next stage should involve sandboxed prospective use in which the framework informs expert review but does not control experimental selection. Only after calibration, multi-site replication, and documented decision benefit should more integrated workflows be considered. Throughout implementation, organizations should define which constraints are mandatory, which trade-offs remain negotiable, how uncertainty is communicated, and who is accountable for escalation. The aim is not to eliminate chemistry judgment, but to expose synthesis assumptions early enough that molecular design and chemical realization can be reasoned about together.

Conclusion

Synthesis-constrained molecular design requires more than adding a synthetic-accessibility score after a generator has optimized a structure. The proposed Route-Embedded Design Viability theory treats each candidate as a coupled molecule–route hypothesis whose value depends on reaction feasibility, route topology, material access, process burden, uncertainty, and the declared operational context. This organization allows synthesis evidence to revise the molecular proposal during generation and preserves trade-offs between molecular promise and route viability without forcing them into a universal scalar

objective. Its output is not proof of synthesis, manufacturability, therapeutic utility, or deployment readiness, but an evidence-structured basis for deciding which molecule–route pairs justify deeper expert and experimental evaluation. Synthesis begins at generation when chemical realization is treated as part of what a molecular design means rather than as a filter applied after the design is complete.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Sánchez-Lengeling B, Aspuru-Guzik A. Inverse molecular design using machine learning: Generative models for matter engineering. Science. 2018;361(6400):360-5. doi:10.1126/science.aat2663
2. Schneider P, Walters WP, Plowright AT, Sieroka N, Listgarten J, Goodnow RA, et al. Rethinking drug design in the artificial intelligence era. Nat Rev Drug Discov. 2020;19(5):353-64. doi:10.1038/s41573-019-0050-3
3. 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
4. Bilodeau C, Jin W, Jaakkola TS, Barzilay R, Jensen KF. Generative models for molecular discovery: Recent advances and challenges. Wiley Interdiscip Rev Comput Mol Sci. 2022;12(5). doi:10.1002/wcms.1608
5. 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
6. Brown N, Fiscato M, Segler MHS, Vaucher AC. GuacaMol: Benchmarking models for de novo molecular design. J Chem Inf Model. 2019;59(3):1096-108. doi:10.1021/acs.jcim.8b00839
7. 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. doi:10.3389/fphar.2020.565644
8. Renz P, Van Rompaey D, Wegner JK, Hochreiter S, Klambauer G. On failure modes in molecule generation and optimization. Drug Discov Today Technol. 2019;32-33:55-63. doi:10.1016/j.ddtec.2020.09.003
9. Langevin M, Vuilleumier R, Bianciotto M. Explaining and avoiding failure modes in goal-directed generation of small molecules. J Cheminform. 2022;14(1):20. doi:10.1186/s13321-022-00601-y
10. Thomas M, O'Boyle NM, Bender A, de Graaf C. MolScore: A scoring, evaluation and benchmarking framework for generative models in de novo drug design. J Cheminform. 2024;16(1):64. doi:10.1186/s13321-024-00861-w
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. doi:10.1021/acs.jcim.7b00622
12. 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
13. Genheden S, Thakkar A, Chadimová V, Reymond JL, Engkvist O, Bjerrum EJ. AiZynthFinder: A fast, robust and flexible open-source software for retrosynthetic planning. J Cheminform. 2020;12(1):70. doi:10.1186/s13321-020-00472-1
14. Li J, Eastgate MD. Making better decisions during synthetic route design: Leveraging prediction to achieve greenness-by-design. React Chem Eng. 2019;4(9):1595-607. doi:10.1039/C9RE00019D
15. Sherer EC, Bagchi A, Kosjek B, Maloney KM, Peng Z, Robaire SA, et al. Driving aspirational process mass intensity using simple structure-based prediction. Org Process Res Dev. 2022;26(5):1405-10. doi:10.1021/acs.oprd.1c00477
16. Button AL, Merk D, Hiss JA, Schneider G. Automated de novo molecular design by hybrid machine intelligence and rule-driven chemical synthesis. Nat Mach Intell. 2019;1(7):307-15. doi:10.1038/s42256-019-0067-7
17. Grisoni F, Huisman BJH, Button AL, Moret M, Atz K, Merk D, et al. Combining generative artificial intelligence and on-chip synthesis for de novo drug design. Sci Adv. 2021;7(24). doi:10.1126/sciadv.abg3338
18. Fialková V, Zhao J, Papadopoulos K, Engkvist O, Bjerrum EJ, Kogej T, et al. LibINVENT: Reaction-based generative scaffold decoration for in silico library design. J Chem Inf Model. 2022;62(9):2046-63. doi:10.1021/acs.jcim.1c00469
19. Parrot M, Tajmouati H, da Silva VBR, Atwood BR, Fourcade R, Gaston-Mathé Y, et al. Integrating synthetic accessibility with AI-based generative drug design. J Cheminform. 2023;15(1):83. doi:10.1186/s13321-023-00742-8
20. Guo J, Schwaller P. Directly optimizing for synthesizability in generative molecular design using retrosynthesis models. Chem Sci. 2025;16(16):6943-56. doi:10.1039/D5SC01476J
21. Winter R, Montanari F, Steffen A, Briem H, Noé F, Clevert DA. Efficient multi-objective molecular optimization in a continuous latent space. Chem Sci. 2019;10(34):8016-24. doi:10.1039/C9SC01928F

22. Liu X, Ye K, van Vlijmen HWT, Emmerich MTM, IJzerman AP, van Westen GJP. DrugEx v2: De novo design of drug molecules by Pareto-based multi-objective reinforcement learning in polypharmacology. J Cheminform. 2021;13(1):85. doi:10.1186/s13321-021-00561-9
23. Guo J, Schwaller P. Augmented Memory: Sample-efficient generative molecular design with reinforcement learning. JACS Au. 2024;4(6):2160-72. doi:10.1021/jacsau.4c00066
24. Yin X, Wang X, Wu Z, Li Q, Kang Y, Deng Y, et al. Syn-MolOpt: A synthesis planning-driven molecular optimization method using data-derived functional reaction templates. J Cheminform. 2025;17(1):27. doi:10.1186/s13321-025-00975-9
25. 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
26. Thakkar A, Kogej T, Reymond JL, Engkvist O, Bjerrum EJ. Datasets and their influence on the development of computer assisted synthesis planning tools in the pharmaceutical domain. Chem Sci. 2020;11(1):154-68. doi:10.1039/C9SC04944D
27. Torren-Peraire P, Hassen AK, Genheden S, Verhoeven J, Clevert DA, Preuss M, et al. Models matter: The impact of single-step retrosynthesis on synthesis planning. Digit Discov. 2024;3(3):558-72. doi:10.1039/D3DD00252G
28. Hastedt F, Bailey RM, Hellgardt K, Yaliraki SN, del Rio-Chanona EA, Zhang D. Investigating the reliability and interpretability of machine learning frameworks for chemical retrosynthesis. Digit Discov. 2024;3(6):1194-212. doi:10.1039/D4DD00007B
29. 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
30. Pasquini M, Stenta M. LinChemIn: SynGraph—a data model and a toolkit to analyze and compare synthetic routes. J Cheminform. 2023;15(1):41. doi:10.1186/s13321-023-00714-y
31. Saigiridharan L, Hassen AK, Lai H, Torren-Peraire P, Engkvist O, Genheden S. AiZynthFinder 4.0: Developments based on learnings from 3 years of industrial application. J Cheminform. 2024;16(1):57. doi:10.1186/s13321-024-00860-x
32. Hassen AK, Šícho M, van Aalst YJ, Huizenga MCW, Reynolds DNR, Luukkonen S, et al. Generate what you can make: Achieving in-house synthesizability with readily available resources in de novo drug design. J Cheminform. 2025;17(1):41. doi:10.1186/s13321-024-00910-4