Pharmacophore

ISSN-2229-5402

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

POTENCY IS ONLY ONE OBJECTIVE IN THE COMPUTATIONAL DESIGN OF MEDICINES WITH VIABLE PHARMACEUTICAL FUTURES

Klaus Weber[1]*, Petra Novakova[2], Li Zhang[3], Stefan Bauer[4]

1. ***Department of Multi-Objective Drug Design, Faculty of Pharmaceutical Sciences, University of Bonn, Bonn, Germany.***
2. ***Department of Pharmaceutical Future Assessment, Faculty of Pharmacy, University of South Bohemia, České Budějovice, Czech Republic.***
3. ***Department of Viable Medicine Design, Faculty of Pharmacy, Zhejiang University, Hangzhou, China.***
4. ***Department of Beyond-Potency Optimization, Faculty of Pharmacy, University of Freiburg, Freiburg, Germany.***

ARTICLE INFO

Received:
24 October 2024
Received in revised form:
17 December 2024
Accepted:
19 December 2024
Available online:
28 December 2024

Keywords: Multiobjective molecular design, Pharmaceutical viability, Potency optimization, Pareto reasoning, Drug developability, Preference specification

ABSTRACT

Computational molecular design increasingly enables the generation, ranking, and optimization of chemical structures against predicted biological and physicochemical properties. However, many design systems continue to privilege potency or a closely related target-activity score as the dominant measure of molecular quality. This emphasis creates a conceptual mismatch between optimization success and pharmaceutical viability because a highly potent structure may remain unsuitable owing to inadequate exposure, undesirable target interactions, safety liabilities, synthetic difficulty, formulation constraints, or weak alignment with the intended therapeutic context. This Original Multiobjective Design Theory Article develops a proposed account of medicine design in which computational optimization is directed toward viable pharmaceutical futures rather than isolated property maxima. The article conceptualizes candidate quality as a position within a multidimensional objective space comprising efficacy, selectivity, exposure, safety, and manufacturability, with each domain represented by imperfect evidence and subject to programme-specific constraints. It further distinguishes compensatory trade-offs from non-compensatory failures, proposes a viable pharmaceutical region within which Pareto reasoning becomes scientifically meaningful, and argues that preferences should change as evidence accumulates across discovery stages. The central contribution is therefore not a universal scoring formula but a theory for organizing objectives, constraints, uncertainty, and preference specification around pharmaceutical decisions. The proposed account remains conceptual: it does not establish validated thresholds, universally correct objective weights, prospective workflow superiority, clinical utility, regulatory acceptability, or deployment readiness. Its principal implication is that computational medicine design should be evaluated by the quality and transparency of the candidate sets and decisions it supports, rather than by potency improvement or aggregate benchmark performance alone.

To Cite This Article: Weber K, Novakova P, Zhang L, Bauer S. Potency Is Only One Objective in the Computational Design of Medicines with Viable Pharmaceutical Futures. Pharmacophore. 2024;15(6):127-38. https://doi.org/10.51847/0J4U9dEvXj

Introduction

Computational drug design has moved beyond the isolated calculation of molecular descriptors toward systems that can generate structures, predict properties, prioritize experiments, and support iterative design–make–test cycles. Automation can accelerate selected search and design activities, but it does not remove the need to define what constitutes a scientifically and pharmaceutically acceptable result. Medicinal chemistry remains a decision process conducted under incomplete evidence, competing objectives, and substantial experimental uncertainty. Consequently, the value of automation depends not merely on how efficiently a system searches chemical space but on whether its search objective reflects the conditions under which a molecule could plausibly progress toward a medicine [1].

Corresponding Author: Klaus Weber; Department of Multi-Objective Drug Design, Faculty of Pharmaceutical Sciences, University of Bonn, Bonn, Germany. E-mail: klaus.weber@uni-bonn.de.

Machine-learning methods now contribute across target identification, molecular property prediction, virtual screening, de novo design, biomarker development, preclinical assessment, and other parts of the discovery and development continuum. These applications differ markedly in their data structures, evidentiary maturity, permissible claims, and consequences of error. A prediction that is useful for ranking early hypotheses may be inadequate for excluding a safety risk or supporting candidate nomination. The relevance of computational output is therefore conditional on the task, the data-generating process, the intended decision, and the availability of experimental confirmation [2]. Yet molecular design systems are still frequently described through a narrow vocabulary of improved scores, optimized properties, or generated candidates without specifying whether those outputs address a coherent pharmaceutical decision.

This problem is particularly consequential when target potency becomes the organizing objective. Potency is experimentally meaningful and often necessary, but it is neither a complete representation of efficacy nor a substitute for exposure, selectivity, safety, or manufacturability. A design process may produce molecules with increasingly favorable predicted activity while simultaneously moving toward regions of chemical space associated with poor permeability, rapid clearance, off-target effects, reactive functionality, synthetic complexity, or formulation difficulty. Reconsidering drug design in the artificial-intelligence era therefore requires the unit of analysis to shift from isolated model output to the pharmaceutical decision that the output is intended to inform [3]. The unresolved theoretical question is how heterogeneous objectives should be represented and related without hiding their conflicts inside a single apparently precise score.

Molecular generation and property-prediction methods provide important enabling capabilities, but neither capability alone establishes that a proposed structure has a viable pharmaceutical future [4]. Moreover, apparently strong computational performance may be misleading when validation data are redundant, endpoints are weak proxies, applicability domains are ignored, or retrospective benchmarks are treated as evidence of prospective usefulness [5]. This article addresses that gap by proposing a multiobjective design theory of viable pharmaceutical futures. Its central construct is the viable pharmaceutical region: a context-dependent portion of objective space in which minimum requirements for efficacy, selectivity, exposure, safety, and manufacturability are jointly satisfied sufficiently to justify continued investigation. Within that region, candidates are not reduced to a universal ranking. Instead, non-dominated alternatives are compared through explicit, stage-sensitive preferences and uncertainty statements. The contribution is an organizing theory rather than an empirically validated model, scoring function, regulatory standard, or operational decision tool.

Why Potency-Centered Optimization Produces Fragile Candidates

A potency-centered candidate is fragile when its apparent superiority depends on treating target activity as the dominant or sufficient representation of pharmaceutical value. Fragility does not mean that potency is unimportant; it means that the candidate's ranking is unstable once additional requirements are introduced. Evidence from pharmaceutical research and development indicates that progression depends on several interacting determinants, including biological understanding of the target, adequate exposure at the relevant tissue, acceptable safety, identification of an appropriate patient population, and alignment with a credible development strategy [6]. A molecular structure can therefore be strong with respect to one experimentally measured property while remaining weak as a development proposition. In the proposed theory, potency is a necessary candidate attribute only when the therapeutic hypothesis requires it, whereas pharmaceutical viability is a relational judgment concerning the molecule, biological system, dosing context, formulation, manufacturing route, and intended use.

The distinction between biochemical potency and therapeutically relevant activity is especially important. A molecule may bind strongly in an isolated assay but fail to achieve adequate unbound concentration at the site of action, may distribute into tissues associated with toxicity, or may interact with unintended targets at clinically relevant exposure. Analyses of clinical development failure have emphasized that optimization centered on potency and nominal specificity can remain misleading when tissue exposure and tissue selectivity are insufficiently examined [7]. Exposure is not a secondary convenience added after potency optimization; it determines whether the concentration required for the intended pharmacology can plausibly be reached and sustained. Likewise, selectivity cannot be interpreted solely from a small assay panel or from nominal affinity ratios without considering free concentrations, metabolites, tissue distribution, target expression, and the therapeutic window. A potency optimum that cannot be exposed safely is not a weaker version of a medicine; it represents a different and potentially non-viable pharmaceutical state.

Potency-centered fragility can also arise when simple physicochemical rules are treated as universal surrogates for developability. Property heuristics can guide early design, but the distributions of approved oral drugs and the chemical strategies used to achieve exposure have changed over time. This evolution limits the interpretation of any fixed rule as a complete boundary between viable and non-viable molecules [8]. The theoretical implication is not that property constraints should be abandoned. Rather, property values must be interpreted in relation to modality, route of administration, target location, dose, formulation strategy, transporter effects, metabolic pathways, and the evidentiary stage of the programme. A candidate may violate a conventional property expectation yet remain developable through a coherent set of structural and pharmaceutical adaptations. Conversely, formal compliance with common property ranges does not establish adequate exposure, safety, or therapeutic value.

Experience with compounds beyond conventional property boundaries further demonstrates that successful development is context-dependent and cannot be inferred from potency alone [9]. High molecular size, lipophilicity, conformational complexity, or other challenging attributes may sometimes be managed, but doing so can require deliberate control of permeability, solubility, efflux, metabolism, formulation, and route feasibility. These dependencies create asymmetry among

objectives: a modest loss of potency may sometimes be acceptable if exposure or selectivity improves, whereas an unresolved genotoxic liability, an unattainable tissue concentration, or the absence of a credible synthetic route may be non-compensatory. Potency-centered optimization becomes fragile precisely because it tends to postpone these asymmetries. The proposed theory therefore treats fragility as an objective-coverage problem: a candidate is insufficiently characterized, or improperly preferred, when its high score in one domain conceals absent, uncertain, or unacceptable evidence in another.

Table 1 organizes the evidence, constructs, and interpretation boundaries needed to develop why potency-centered optimization produces fragile candidates within the article's central argument.

Table 1. Evidence domains, core questions, scientific requirements, and interpretation boundaries for Why potency-centered optimization produces fragile candidates

Construct or evidence domain	Core question	Scientific basis	Role in the article	Failure or overclaiming risk	Boundary statement
Potency	Does the molecule produce the required target-level pharmacology under an appropriate assay context?	Biochemical and cellular activity establish one dimension of molecular action but do not establish exposure, selectivity, safety, or therapeutic effect.	Defines an important but incomplete optimization objective.	Treating a low activity value or high predicted affinity as equivalent to efficacy.	Potency is neither sufficient evidence of therapeutic benefit nor a universal primary objective.
Target and biological rationale	Is the intended target relationship sufficiently supported for the proposed intervention?	Candidate value depends on the validity, direction, and context of the biological hypothesis.	Connects molecular activity to the intended therapeutic mechanism.	Assuming that target engagement automatically produces the desired disease-level outcome.	Computational association or binding prediction does not establish mechanism, causality, or clinical relevance.
Tissue exposure	Can an active unbound concentration plausibly be reached and maintained at the relevant site?	Absorption, distribution, clearance, protein binding, transport, metabolism, and dosing jointly influence effective exposure.	Prevents potency from being interpreted without pharmacokinetic context.	Comparing potency with total plasma concentration or ignoring tissue-specific distribution.	Predicted exposure requires experimental confirmation and does not alone establish efficacy.
Selectivity and polypharmacology	Are intended and unintended target interactions compatible with the therapeutic hypothesis?	Biological effects may arise from selective action, rational multitarget activity, or undesirable off-target engagement.	Distinguishes deliberate pharmacological breadth from uncontrolled promiscuity.	Equating narrow screening-panel selectivity with in vivo safety or mechanism specificity.	Selectivity is programme-specific and must be assessed at relevant exposure and biological context.
Safety	Are observed or predicted liabilities compatible with the intended dose, duration, population, and therapeutic need?	Safety emerges from on-target effects, off-target interactions, metabolites, reactive chemistry, organ exposure, and patient context.	Establishes a major non-compensatory boundary within the viable pharmaceutical region.	Allowing high potency or an aggregate score to offset an unacceptable liability.	In silico alerts organize risk; they do not exclude toxicity or define an acceptable therapeutic margin.
Physicochemical developability	Are molecular properties compatible with the intended route, formulation, exposure, and dose?	Solubility, permeability, ionization, lipophilicity, stability, and solid-state behavior interact rather than acting as isolated pass–fail rules.	Replaces universal property cutoffs with context-dependent interpretation.	Treating heuristic compliance as proof of developability or heuristic violation as automatic failure.	Property rules are decision aids, not validated universal definitions of pharmaceutical viability.
Synthetic and manufacturing feasibility	Is there a credible path from proposed structure to reproducible material of suitable quality?	Route availability, starting materials, stereochemical control, purification, scalability, solid form, and formulation affect progression.	Extends molecular design from structural generation toward a viable pharmaceutical future.	Equating a retrosynthetic route, accessibility score, or purchasable precursor with manufacturability.	Computational feasibility is preliminary and requires laboratory, process, analytical, and formulation evidence.
Evidence completeness and uncertainty	Which objective values are measured, predicted, missing, or outside their applicability domain?	Different evidence sources have unequal reliability, transferability, and decision relevance.	Makes candidate fragility visible when a strong score depends on absent or uncertain domains.	Reporting a complete-looking profile assembled from weak, correlated, or uncalibrated predictors.	Missing evidence is not favorable evidence, and model confidence is not calibrated uncertainty.

The Multidimensional Pharmaceutical Objective Space

The multidimensional pharmaceutical objective space begins with the recognition that medicines act through heterogeneous biological targets embedded in complex physiological systems. The known therapeutic target landscape includes proteins, nucleic acids, receptors, enzymes, channels, transporters, and other target classes with different degrees of biological validation, tractability, tissue distribution, and disease relevance [10]. Target engagement is therefore foundational but not equivalent to therapeutic viability. Even a well-characterized interaction must be interpreted through the direction and duration

of modulation, disease biology, compensatory pathways, patient heterogeneity, and the relationship between molecular exposure and functional response. In the proposed theory, efficacy is consequently broader than potency: it represents the plausibility that a molecule can produce the required pharmacological effect under a relevant exposure and biological context. Potency contributes to this objective, but it cannot stand in for the complete efficacy claim.

Selectivity occupies a second objective domain, although its direction is not universally "more is better." Some programmes require narrow target selectivity to protect a therapeutic window or clarify mechanism, whereas others may benefit from deliberately balanced activity across multiple targets. Medicinal-chemistry accounts of designed polypharmacology show that multitarget action can be purposeful rather than merely undesirable promiscuity [11]. The relevant objective is therefore not maximal selectivity in the abstract, but an exposure-aware interaction profile aligned with the therapeutic hypothesis. This profile should distinguish intended targets, tolerated secondary pharmacology, uncertain interactions, and unacceptable off-target activity. Preference specification becomes necessary because two molecules may be non-dominated: one may provide cleaner target selectivity, while another offers a rational multitarget pattern with stronger expected pathway coverage. The theory does not resolve that choice automatically; it requires the preference and its biological justification to be made explicit.

Safety forms a distinct but interconnected objective because unwanted molecular activity may alter both candidate risk and interpretation of mechanism. Experimental reassessment of oncology compounds has demonstrated that off-target toxicity can be a common basis of apparent drug action, including cases in which the presumed target mechanism does not adequately explain the observed effect [12]. This evidence supports two boundaries within the multidimensional space. First, computational or phenotypic association should not be reported as target-specific mechanism without orthogonal validation. Second, a desirable effect cannot be evaluated independently from the exposure at which unintended interactions occur. Safety is also only partly compensatory. Some tolerability trade-offs may depend on indication, duration, patient population, monitoring, and unmet need, but serious or mechanistically credible liabilities may define exclusion constraints regardless of predicted potency. The viable pharmaceutical region must therefore contain both graded preferences and hard boundaries.

Exposure describes whether pharmacologically relevant concentrations can plausibly occur at the intended site for an appropriate duration while remaining compatible with safety. In silico ADME and pharmacokinetic methods can support compound prioritization, experimental design, and interpretation, but their utility depends on appropriate data, endpoints, assumptions, and integration with measured evidence [13]. Exposure cannot be represented adequately by a single permeability, clearance, or bioavailability prediction. It emerges from interacting processes, including solubility, absorption, distribution, protein binding, metabolism, transport, route, dose, formulation, and active metabolites. The fifth domain, manufacturability, concerns whether the molecule can become reproducible pharmaceutical material through a credible synthetic, process, analytical, solid-state, and formulation path. Together, efficacy, selectivity, exposure, safety, and manufacturability define the proposed objective space. They are analytically distinguishable but causally and practically coupled; no domain score is sufficient alone, and the acceptable region remains programme-specific, evidence-dependent, and subject to future experimental validation.

Formalizing Trade-Offs Among Efficacy, Selectivity, Exposure, Safety, and Manufacturability

A formal account of multiobjective medicine design must distinguish objective values from the evidence used to estimate them. Let a candidate molecule be represented by (x), and let its provisional pharmaceutical profile be described by objective functions for efficacy, selectivity, exposure, safety, and manufacturability. These functions are not direct measurements of eventual therapeutic value. Each may instead combine assay results, computational predictions, expert interpretations, and unresolved evidence gaps. Multiparameter molecular optimization demonstrates that properties with orthogonal or conflicting tendencies can be encoded within a common computational search problem [14]. The proposed theory extends that computational idea by requiring every objective to be accompanied by its biological meaning, evidence provenance, uncertainty, applicable context, and conditions under which compensation by another objective is impermissible.

Formalization should therefore begin with three distinct elements: objectives, constraints, and uncertainties. Objectives describe directions of improvement, such as increasing relevant target engagement, improving tissue exposure, reducing undesirable interactions, or simplifying synthesis. Constraints specify requirements that a candidate must satisfy before trade-off comparison is legitimate. Uncertainty describes the reliability and applicability of each estimated value. Continuous latent-space optimization can search several predicted molecular properties efficiently, but the resulting candidates inherit the limitations of the learned representation, training distribution, and property predictors [15]. A numerically favorable position in latent space consequently does not establish that a candidate occupies a feasible pharmaceutical region. The formal representation must preserve the distinction between optimizing a prediction and improving the underlying pharmaceutical property.

The relationship among objectives cannot be assumed to be fully compensatory. A weighted sum permits a large improvement in one dimension to offset deterioration in another, even when the latter represents an unacceptable liability. Such compensation may be reasonable for selected graded preferences, such as balancing moderate potency against improved solubility, but not for every failure. Generative models have produced molecules with attractive computational scores that are nevertheless difficult to synthesize, illustrating how optimizing a primary objective can shift proposals away from practical chemical feasibility [16]. In the proposed theory, each programme must therefore identify non-compensatory constraints, conditionally compensatory objectives, and preference-sensitive objectives. This classification prevents an excellent potency

prediction from mathematically cancelling an unresolved safety alert, an inaccessible route, or an exposure requirement that cannot plausibly be met.

Manufacturability requires particular care because synthetic accessibility scores compress several different questions into a convenient proxy. Retrosynthesis-informed classifiers can provide rapid estimates of whether a molecular proposal resembles structures for which a planning system can identify plausible routes [17]. However, route identification does not establish reagent availability, stereochemical control, yield, impurity management, purification feasibility, solid-form behavior, scale-up robustness, formulation compatibility, or supply continuity. The proposed formalization therefore treats synthesizability as an early feasibility signal nested within the broader manufacturability objective. Candidate comparison should record which feasibility claims are calculated, planner-derived, experimentally demonstrated, or not yet assessed. **Table 2** organizes the evidence, constructs, and interpretation boundaries needed to develop formalizing trade-offs among efficacy, selectivity, exposure, safety, and manufacturability within the article's central argument.

Table 2. Components, relationships, uncertainties, and validation needs within Formalizing trade-offs among efficacy, selectivity, exposure, safety, and manufacturability

Component or layer	Inputs	Core function	Expected output	Uncertainty or limitation	Validation requirement
Objective-definition layer	Therapeutic hypothesis, target product profile, assay endpoints, route, population, and programme stage	Defines what efficacy, selectivity, exposure, safety, and manufacturability mean for the specific programme	A programme-specific objective set with stated directionality	Definitions may rely on incomplete biology or inappropriate proxy endpoints	Cross-functional review and revision as biological and pharmaceutical evidence accumulates
Evidence layer	Experimental measurements, computational predictions, literature evidence, expert judgment, and missing-data indicators	Separates measured evidence from predicted or inferred values	Evidence-tagged candidate profiles	Evidence sources may differ in scale, reliability, transferability, and applicability domain	Orthogonal assays, repeated measurements, provenance documentation, and context checks
Constraint layer	Minimum activity, exposure feasibility, prohibited liabilities, route requirements, and material-quality conditions	Excludes candidates that fail non-compensatory requirements	A provisional feasible pharmaceutical region	Thresholds may be uncertain, stage-dependent, or programme-specific	Prospective justification of constraints and sensitivity analysis around uncertain thresholds
Trade-off layer	Candidate objective profiles within the feasible region	Identifies improving, worsening, and conflicting objective relations	Non-dominated alternatives and explicit trade-off descriptions	Correlated predictors or incompatible units can create misleading comparisons	Comparison with experimentally updated profiles and alternative formalizations
Uncertainty layer	Prediction intervals, assay variability, model applicability, route confidence, and missing evidence	Qualifies objective values and dominance relations	Uncertainty-aware candidate profiles and unstable-comparison flags	Model confidence may be mistaken for calibrated uncertainty	Calibration studies, repeated assays, external testing, and uncertainty propagation
Compensability layer	Safety boundaries, exposure requirements, therapeutic need, formulation options, and programme preferences	Distinguishes permissible trade-offs from inadmissible compensation	Hard constraints, conditional trade-offs, and preference-sensitive dimensions	Normative choices may be hidden inside mathematical weights	Explicit multidisciplinary approval and documented rationale
Manufacturability layer	Retrosynthetic plans, starting materials, process knowledge, analytical data, solid-state evidence, and formulation considerations	Extends synthetic feasibility toward a credible material and manufacturing path	Graded evidence of synthesis and pharmaceutical-material feasibility	A route-found status can overstate laboratory or process readiness	Executed synthesis, process evaluation, impurity control, solid-form studies, and formulation testing
Decision layer	Feasible alternatives, uncertainty, programme priorities, resource limits, and human judgment	Supports selection, retention, rejection, or further experimentation	A traceable decision with unresolved risks stated	A formal model cannot determine the ethically or strategically correct preference alone	Prospective decision audit and comparison with subsequent experimental evidence

Pareto Reasoning and Preference Specification in Medicine Design

Pareto reasoning is useful when improvement in one pharmaceutical objective cannot be obtained without deterioration in another. A candidate is non-dominated when no available alternative is at least as favorable across all relevant objectives and strictly more favorable in one. This definition prevents premature collapse of efficacy, selectivity, exposure, safety, and manufacturability into a single universal score. Computer-aided multiobjective optimization can expose non-dominated alternatives and make trade-offs visible before a preference is imposed [18]. The proposed viable pharmaceutical region adds

an essential qualification: Pareto comparison should occur only among candidates that satisfy programme-specific minimum constraints. A molecule with an unacceptable safety liability should not remain a preferred Pareto solution merely because its potency and synthetic accessibility are exceptional.

Feasibility must therefore precede preference. Unconstrained optimization can exploit unsupported areas of a molecular representation or generate structures that formally improve predicted objectives while departing from chemically valid or data-supported regions. Constrained Bayesian optimization illustrates how feasibility conditions can be incorporated into molecular search rather than treated solely as a post-generation filter [19]. Pharmaceutical feasibility is broader than chemical validity, however. It may include minimum exposure plausibility, restrictions on reactive functionality, essential selectivity conditions, route credibility, or formulation requirements. The proposed theory treats these conditions as revisable scientific commitments rather than universal constants. Their purpose is to prevent preferences from legitimizing candidates that should not yet enter the trade-off set.

Preference specification begins after feasible non-dominated alternatives have been identified. Conditional molecular generation can steer proposed structures toward selected descriptor values, demonstrating that design requirements can be represented directly within generative processes [20]. Nevertheless, satisfaction of descriptor conditions remains a proxy achievement. A requested lipophilicity range does not prove oral exposure, a predicted selectivity score does not establish an acceptable safety margin, and a synthetic-accessibility estimate does not demonstrate scalable manufacture. Preferences should consequently be expressed in pharmaceutical language before they are encoded computationally. The programme should state which improvements are valued, which sacrifices are tolerable, which constraints are non-negotiable, and which evidence is too uncertain to support a decisive comparison.

Reward-based molecular design makes the consequences of hidden preferences especially clear. Reinforcement-learning systems optimize the reward they are given and may exploit imperfect reward functions rather than advance the broader therapeutic purpose intended by their designers [21]. The proposed theory therefore requires preference provenance: every weight, priority, constraint, or ordering rule should be linked to an accountable scientific or strategic rationale. Sensitivity analysis should determine whether the preferred candidate changes under plausible alternative preferences or uncertainty assumptions. Stable preference does not prove pharmaceutical success, but unstable preference reveals that candidate selection depends heavily on unresolved judgments. **Table 3** organizes the evidence, constructs, and interpretation boundaries needed to develop pareto reasoning and preference specification in medicine design within the article's central argument.

Table 3. Components, relationships, uncertainties, and validation needs within Pareto reasoning and preference specification in medicine design

Evaluation dimension	What must be tested	Suitable evidence or assessment	Failure signal	Interpretive limitation	Decision relevance
Feasibility integrity	Whether all candidates entering Pareto comparison satisfy the stated minimum constraints	Constraint audit, applicability-domain review, chemistry checks, and experimental evidence status	A non-dominated candidate remains despite failing a non-compensatory requirement	Feasibility boundaries may change as evidence matures	Determines whether trade-off comparison is scientifically permissible
Dominance validity	Whether apparent superiority persists when uncertainty and measurement variability are considered	Uncertainty-aware dominance analysis and repeated or orthogonal measurements	Candidate ordering reverses within plausible uncertainty ranges	Non-dominance does not establish therapeutic adequacy	Identifies alternatives that should remain under consideration
Frontier stability	Whether the non-dominated set is robust to model, data, and objective-definition changes	Alternative predictors, resampling, temporal updating, and sensitivity analysis	Large frontier changes after minor modeling choices	Stability is conditional on the represented candidate set	Indicates whether a trade-off surface is decision-worthy
Preference transparency	Whether weights, ordering rules, and constraints have explicit pharmaceutical rationales	Preference documentation and multidisciplinary review	Preferences are selected solely for algorithmic convenience	Transparent preferences may still be scientifically wrong	Makes normative and strategic judgments contestable
Preference sensitivity	Whether the preferred subset changes under plausible alternative priorities	Weight perturbation, scenario analysis, and outranking comparison	A small preference change produces a different winner	Sensitivity does not identify which preference is correct	Shows whether a single-candidate decision is premature
Proxy adequacy	Whether optimized descriptors represent the intended pharmaceutical objectives	Endpoint validation, mechanistic interpretation, and comparison with higher-fidelity evidence	Descriptor satisfaction fails to translate to the intended property	A validated proxy in one context may not transfer to another	Prevents proxy optimization from being presented as therapeutic progress
Diversity of alternatives	Whether the frontier contains chemically and mechanistically distinct options	Scaffold, series, route, property, and mechanism diversity assessment	Many apparently distinct candidates belong to one narrow solution family	Diversity alone does not imply quality	Preserves fallback options against later-stage failure

Decision traceability	Whether exclusions, retained alternatives, overrides, and unresolved risks are recorded	Versioned decision record and prospective review	Candidate selection cannot be reconstructed or justified	Documentation does not guarantee sound judgment	Enables learning across design–make–test cycles

Figure 1 presents the multiobjective medicine design tension field, showing how the article's key components and boundaries are connected within pareto reasoning and preference specification in medicine design.

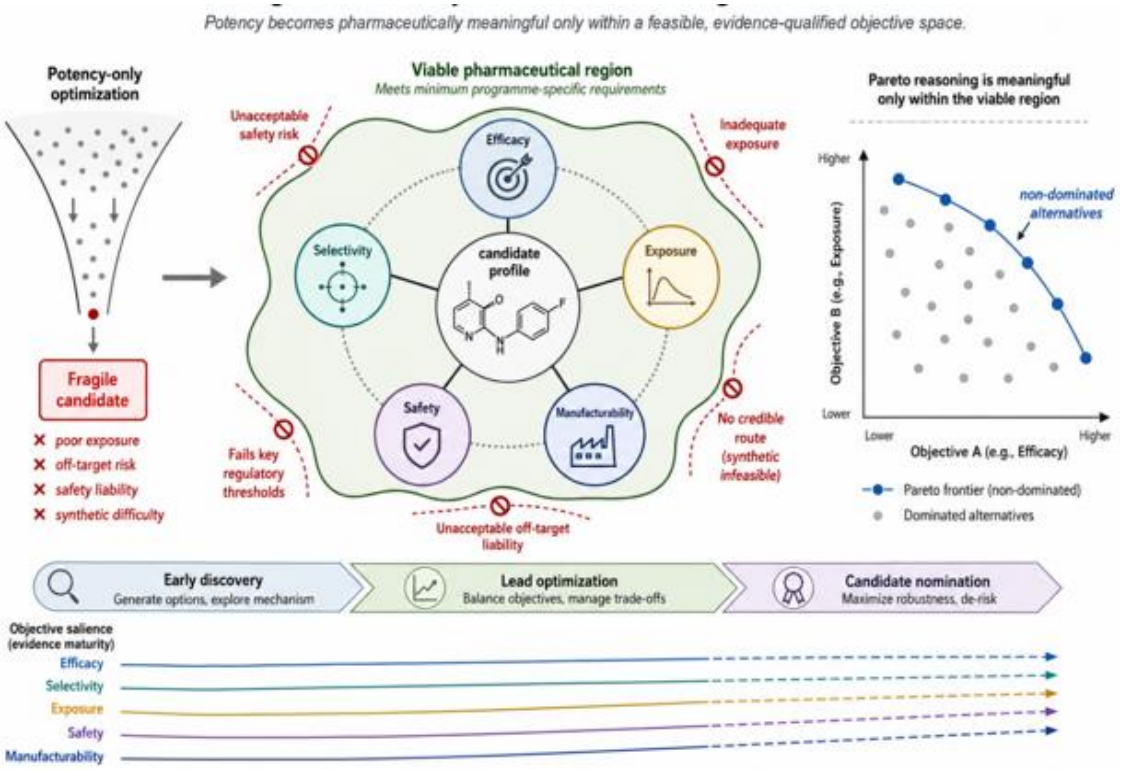

Figure 1. Multiobjective Medicine Design Tension Field.

The figure is an original conceptual synthesis that organizes the article's central contribution across potency-centered optimization produces fragile candidates, multidimensional pharmaceutical objective space, formalizing trade-offs among efficacy, selectivity, exposure, safety, and manufacturability, Pareto reasoning and preference specification in medicine design, dynamic objective weighting across discovery stages, and evaluation principles for multiobjective molecular design systems. 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.

Dynamic Objective Weighting Across Discovery Stages

The relative importance of pharmaceutical objectives changes as a discovery programme moves from early hypothesis generation toward candidate nomination and development preparation. Early stages may tolerate substantial uncertainty while emphasizing tractable activity, chemical diversity, and rapid learning. Later stages require stronger evidence for exposure, selectivity, safety, route robustness, material properties, and dose feasibility. Contemporary medicinal-chemistry optimization increasingly integrates biological, physicochemical, pharmacokinetic, safety, structural, and synthetic information rather than following a fixed sequence dominated by potency [22]. The proposed theory describes this change as dynamic objective weighting, but weighting should not be interpreted only as adjustment of numerical coefficients. It also includes changing constraints, evidence standards, acceptable uncertainty, and the decision consequences attached to failure.

Dynamic weighting should be evidence-responsive rather than automatically time-responsive. An iterative multiobjective system may update generated candidates using model-based assessments of molecular quality, but those assessments remain dependent on the labels, predictors, and reward definitions supplied to the model [23]. A stage transition should therefore occur because the decision has changed and relevant evidence has accumulated, not merely because a workflow has completed a predetermined number of optimization cycles. For example, an early programme may prioritize learning across diverse chemotypes, whereas a later programme may impose stringent exposure and safety constraints. The rationale, evidence trigger, and effect of each change should be versioned so that objective drift is distinguishable from scientifically justified adaptation. Conditional graph generation demonstrates that several molecular requirements can be imposed simultaneously, but it does not establish a universally correct hierarchy among them [24]. The relative value of potency, selectivity, exposure, safety, and manufacturability depends on target biology, therapeutic modality, route, indication, dose, treatment duration, patient

population, available alternatives, and programme strategy. Dynamic preference specification should consequently be anchored to an evolving target product profile and a current decision question. It should also retain prior candidate alternatives when uncertainty is high. Reweighting that silently eliminates previously promising series may reduce portfolio resilience and obscure whether an apparent improvement arose from new evidence or from a changed objective definition.

Diversity should remain an explicit objective throughout the discovery process because aggressive optimization can collapse a candidate population onto a narrow chemical series. Memory-assisted reinforcement learning has been used to encourage exploration of diverse molecular solutions rather than repeatedly generating close analogues around one reward maximum [25]. In pharmaceutical terms, diversity has value when it preserves alternatives with different scaffolds, routes, selectivity patterns, intellectual-property positions, or failure risks. It is not an end in itself, and chemically diverse but uniformly poor candidates do not constitute progress. The proposed theory therefore treats diversity as a risk-management objective whose salience may increase when evidence is uncertain, when a lead series has unresolved liabilities, or when late failure would leave no viable fallback.

Evaluation Principles for Multiobjective Molecular Design Systems

Evaluation must begin by asking whether a system optimizes the intended pharmaceutical objectives rather than whether it merely performs well on convenient computational tasks. Benchmark suites such as GuacaMol enable standardized comparison of molecular generators across distribution-learning and goal-directed tasks [26]. Such benchmarks are valuable for testing whether an algorithm can generate valid structures, reproduce selected distributions, or optimize formal objectives under controlled conditions. They do not establish that the objectives correspond to a credible therapeutic programme, that generated molecules can be synthesized, that predicted properties generalize prospectively, or that the system improves pharmaceutical decisions. Benchmark performance should therefore be interpreted as model-level evidence rather than candidate-level or workflow-level validation.

Generative-model evaluation also commonly examines validity, novelty, uniqueness, and diversity. Platforms such as MOSES make these properties more comparable across methods and help identify failures such as invalid generation or reproduction of training examples [27]. Nevertheless, a novel and diverse molecular set may remain irrelevant to the intended target, outside the applicability domains of property models, synthetically inaccessible, or incompatible with exposure and safety requirements. Multiobjective evaluation should therefore include objective validity, constraint satisfaction, uncertainty calibration, frontier quality, preference sensitivity, chemical and route diversity, and evidence progression. No single metric should be interpreted as a sufficient measure of pharmaceutical usefulness.

Evaluation design must also make data provenance, preprocessing, task construction, split strategy, model selection, baselines, uncertainty handling, and reproducibility visible. Best-practice guidance for machine learning in chemistry emphasizes that credible evaluation depends on transparent data and workflows, appropriate validation, and recognition of domain-specific limitations [28]. For multiobjective design, this requirement extends to every predictor contributing to a candidate profile. A system may appear to integrate five objectives while relying on several correlated models trained on overlapping or weakly representative data. Evaluation should identify such dependencies, test calibration within relevant chemical regions, record missing evidence, and determine whether apparent Pareto improvements survive higher-fidelity or experimental assessment.

Generalization is particularly important because random or highly redundant validation structures can reward memorization of familiar chemical series. Ligand-based benchmarks have been shown to provide optimistic assessments when training and test sets share exploitable chemical similarity [29]. Multiobjective systems magnify this problem because several biased predictors may be combined into a convincing but fragile candidate score or frontier. Evaluation should therefore use split designs aligned with the intended claim, including scaffold-, series-, cluster-, or time-aware validation where appropriate, followed by prospective assessment in repeated design–make–test cycles. **Table 4** organizes the evidence, constructs, and interpretation boundaries needed to develop evaluation principles for multiobjective molecular design systems within the article's central argument.

Table 4. Evaluation, implementation, and research priorities arising from Potency Is Only One Objective in the Computational Design of Medicines with Viable Pharmaceutical Futures

Research or implementation priority	Unresolved gap	Required methodological work	Evidence needed	Responsible actors	Expected contribution
Pharmaceutical objective ontology	Objective labels such as efficacy, exposure, safety, and synthesizability are used inconsistently	Develop versioned definitions distinguishing objectives, proxies, assays, predictions, and decision claims	Cross-programme endpoint mappings, assay context, and expert consensus	Medicinal chemists, pharmacologists, DMPK scientists, toxicologists, CMC experts, and data scientists	More comparable objective specification and fewer semantic mismatches
Non-compensatory constraint methods	Aggregate scores can conceal unacceptable liabilities	Compare constrained Pareto, lexicographic, outranking, and hard-boundary approaches	Historical failure analysis and prospective candidate-review evidence	Discovery teams, safety scientists, statisticians, and decision scientists	More transparent exclusion of inadmissible candidates

Uncertainty-aware frontier analysis	Pareto sets are often derived from point predictions	Propagate assay, model, route, and missing-data uncertainty into dominance analysis	Calibrated predictions, repeated assays, and higher-fidelity follow-up	Model developers, experimental scientists, and quantitative pharmacologists	More stable and honest representation of trade-offs
Stage-adaptive preference governance	Objective changes and human overrides are poorly documented	Develop auditable rules for revising constraints, evidence standards, and preferences	Longitudinal programme records and target product profiles	Programme leaders and cross-functional project teams	Clearer alignment between optimization and changing pharmaceutical decisions
Synthesizability-to-manufacturability bridge	Route scores omit process, material, supply, and formulation realities	Integrate retrosynthesis with executed-route, process, solid-state, analytical, and formulation evidence	Laboratory synthesis, scale-up observations, impurity profiles, and material characterization	Synthetic, process, analytical, formulation, and supply specialists	Earlier identification of candidates with credible pharmaceutical futures
Realistic validation and data lineage	Available chemical data may be large but unsuitable for the intended claim	Apply scaffold-, series-, cluster-, temporal-, and context-aware validation with complete lineage	Curated project histories and prospective external compounds	Data stewards, chemoinformaticians, statisticians, and programme scientists	More credible estimates of generalization and applicability
Prospective workflow evaluation	Retrospective benchmarks do not establish design utility	Compare multiobjective support with appropriate alternative workflows across repeated cycles	Predefined recommendations, experimental follow-up, decision records, and failure analysis	Multidisciplinary discovery teams and independent evaluators	Evidence about decision usefulness rather than model performance alone
Reporting and oversight standards	Objectives, weights, exclusions, and overrides can remain hidden	Create reporting requirements for provenance, constraints, uncertainty, preference changes, and human decisions	Multi-site implementation studies and complete audit trails	Journals, research organizations, professional societies, and funders	Reproducible and reviewable multiobjective design practice

Figure 2 presents the evidence, validation, and translation boundary observatory for potency is only one objective in the computational design of medicines with viable pharmaceutical futures, showing how the article's key components and boundaries are connected within evaluation principles for multiobjective molecular design systems.

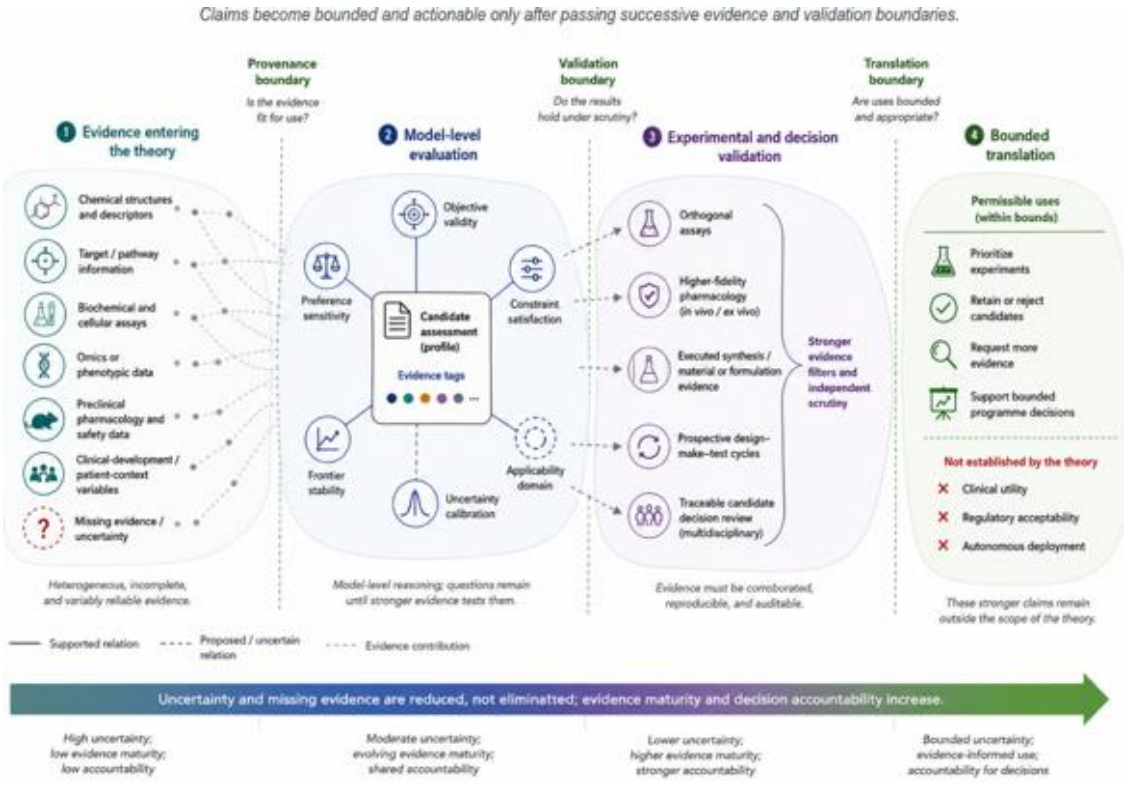

Figure 2. Evidence, Validation, and Translation Boundaries.

The figure is an original conceptual synthesis that shows how claims move from conceptual plausibility through evaluation, uncertainty assessment, and bounded pharmaceutical use. 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.
Include clearly differentiated elements for potency-centered fragility, the multidimensional objective space, formalized trade-offs, Pareto reasoning, preference specification, dynamic objective weighting, and evaluation principles. Use solid connectors only where the manuscript supports the relationship and dashed connectors for proposed, context-dependent, or uncertain relations. Use a clean white background, precise alignment, professional sans-serif typography, thin uniform connectors, generous white space, and a restrained colour-blind-safe palette. Do not make the figure resemble a dashboard, engineering control system, rectangular-box pipeline, circular wheel, marketing infographic, or decorative molecular collage. Do not use logos, interface screenshots, stock laboratory photography, proprietary molecular structures, copied figures, or invented quantitative outputs. Ensure all labels remain legible at journal-column width and that the figure is suitable for high-resolution export.

Limitations and Boundary Conditions

The proposed multiobjective design theory is limited first by the quality and suitability of the evidence available for each objective. Efficacy, selectivity, exposure, safety, and manufacturability are not observed with equal accuracy or at the same stage. Early candidate profiles will often combine measurements with predictions, categorical alerts, expert judgments, and missing values. A formally complete objective vector can therefore create false coherence when its components have different uncertainty, provenance, and relevance. Transparent machine-learning practice can reduce this risk through explicit documentation and appropriate validation, but it cannot convert an inadequately measured pharmaceutical objective into reliable evidence [28]. The theory organizes evidence and decisions; it does not supply missing biology, assays, pharmacokinetic studies, toxicology, executed syntheses, or material characterization.
The viable pharmaceutical region is also context-dependent. Its constraints and preferences may vary with therapeutic target, modality, route, indication, duration, patient population, acceptable dose, unmet need, formulation options, and programme resources. It should not be treated as a universal developability checklist or a fixed mathematical region transferable across programmes. Similarly, Pareto non-dominance does not establish that any candidate is good enough to progress. A frontier may be composed entirely of weak alternatives, may be unstable under uncertainty, or may reflect biased predictors. Validation designs that reward chemical memorization rather than generalization can make apparently sophisticated multiobjective comparisons unreliable [29]. Prospective experimental confirmation remains necessary for claims about candidate quality or workflow value.
A third limitation concerns implementation and misuse. Preference specification introduces scientific, strategic, and sometimes normative judgments that cannot be resolved solely by an optimization algorithm. Human oversight is required to determine which constraints are non-compensatory, which trade-offs are acceptable, and when evidence is sufficient to influence a decision. Synthetic-accessibility predictions, for example, may help prioritize structures but cannot establish manufacturability without executed routes and process evidence [17]. The theory does not establish clinical utility, causal mechanism, regulatory acceptability, commercial value, or deployment readiness. It should not be used to replace candidate-selection governance, safety review, CMC assessment, experimental validation, or formal regulatory requirements.

Research Agenda and Implementation Priorities

The first research priority is to develop clearer pharmaceutical objective definitions and prospective evidence chains. Studies should distinguish the therapeutic property of interest from the computational proxy used to represent it and should document how that proxy is expected to influence a specific decision. Multiobjective analyses should compare hard constraints, lexicographic priorities, constrained Pareto methods, and alternative preference models rather than assuming that a weighted sum is adequate. Because Pareto optimization exposes but does not resolve trade-offs, future work should evaluate how different preference-elicitation procedures affect retained candidates and downstream experimental choices [18]. These studies should include sensitivity to missing data, assay variability, model uncertainty, and disagreement among scientific functions.
A second priority is longitudinal evaluation across discovery stages. Research should examine when objective definitions, constraints, and evidence standards should change and whether those changes improve continuity between hit discovery, lead optimization, candidate nomination, and predevelopment. The changing landscape of medicinal-chemistry optimization suggests that property priorities are increasingly integrated and stage-dependent [22]. Prospective studies should therefore record objective revisions, human overrides, retained fallback series, and the reasons candidates are rejected or advanced. Such work would test whether dynamic weighting improves learning and risk management without prematurely narrowing chemical diversity or allowing early surrogate goals to persist after they have lost decision relevance.
The third priority is infrastructure for realistic, reproducible, and multidisciplinary evaluation. This includes versioned datasets with lineage, context-aware validation splits, calibrated uncertainty, standardized reporting of objectives and constraints, links between retrosynthesis and executed manufacturing evidence, and auditable decision records. Benchmark platforms remain useful for controlled algorithmic comparison, but they should be connected to higher levels of evidence rather than interpreted as substitutes for prospective pharmaceutical assessment [26]. Implementation should proceed in stages: retrospective audit,

silent prospective testing, bounded decision support with human review, and only then broader use if justified by accumulated evidence. No stage should be interpreted as automatic authorization for clinical, regulatory, or autonomous decision-making.

Conclusion

Potency is an indispensable molecular property in many discovery programmes, but it is only one element of a candidate's pharmaceutical future. This article has proposed a multiobjective design theory in which efficacy, selectivity, exposure, safety, and manufacturability are represented as distinct but interacting objectives; minimum constraints define a provisional viable pharmaceutical region; Pareto reasoning preserves meaningful alternatives; explicit preferences determine which trade-offs are acceptable; and objective salience changes as evidence and decisions mature. The contribution is intended to make hidden assumptions, non-compensatory failures, uncertainty, and evidence boundaries visible rather than to prescribe a universal score or permanent optimum. Its usefulness will depend on programme-specific definitions, credible data, experimental confirmation, multidisciplinary judgment, transparent evaluation, and prospective testing. Computational design becomes pharmaceutically meaningful not when it produces the molecule with the highest isolated score, but when it supports an accountable decision among candidates whose combined evidence justifies continued investigation.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Schneider G. Automating drug discovery. Nat Rev Drug Discov. 2018;17(2):97-113. doi:10.1038/nrd.2017.232
2. Vamathevan J, Clark D, Czodrowski P, Dunham I, Ferran E, Lee G, et al. Applications of machine learning in drug discovery and development. Nat Rev Drug Discov. 2019;18(6):463-77. doi:10.1038/s41573-019-0024-5
3. Schneider P, Walters WP, Plowright AT, Sieroka N, Listgarten J, Goodnow RA Jr, 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
4. Walters WP, Barzilay R. Applications of deep learning in molecule generation and molecular property prediction. Acc Chem Res. 2021;54(2):263-70. doi:10.1021/acs.accounts.0c00699
5. 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
6. Morgan P, Brown DG, Lennard S, Anderton MJ, Barrett JC, Eriksson U, et al. Impact of a five-dimensional framework on R&D productivity at AstraZeneca. Nat Rev Drug Discov. 2018;17(3):167-81. doi:10.1038/nrd.2017.244
7. Sun D, Gao W, Hu H, Zhou S. Why 90% of clinical drug development fails and how to improve it? Acta Pharm Sin B. 2022;12(7):3049-62. doi:10.1016/j.apsb.2022.02.002
8. Shultz MD. Two decades under the influence of the Rule of Five and the changing properties of approved oral drugs. J Med Chem. 2019;62(4):1701-14. doi:10.1021/acs.jmedchem.8b00686
9. DeGoey DA, Chen HJ, Cox PB, Wendt MD. Beyond the Rule of 5: Lessons learned from AbbVie's drugs and compound collection. J Med Chem. 2018;61(7):2636-51. doi:10.1021/acs.jmedchem.7b00717
10. Santos R, Ursu O, Gaulton A, Bento AP, Donadi RS, Bologa CG, et al. A comprehensive map of molecular drug targets. Nat Rev Drug Discov. 2017;16(1):19-34. doi:10.1038/nrd.2016.230
11. Proschak E, Stark H, Merk D. Polypharmacology by design: A medicinal chemist's perspective on multitargeting compounds. J Med Chem. 2019;62(2):420-44. doi:10.1021/acs.jmedchem.8b00760
12. Lin A, Giuliano CJ, Palladino A, John KM, Abramowicz C, Yuan ML, et al. Off-target toxicity is a common mechanism of action of cancer drugs undergoing clinical trials. Sci Transl Med. 2019;11(509). doi:10.1126/scitranslmed.aaw8412
13. Lombardo F, Desai PV, Arimoto R, Desino KE, Fischer H, Keefer CE, et al. In silico absorption, distribution, metabolism, excretion, and pharmacokinetics (ADME-PK): Utility and best practices. An industry perspective from the International Consortium for Innovation through Quality in Pharmaceutical Development. J Med Chem. 2017;60(22):9097-113. doi:10.1021/acs.jmedchem.7b00487
14. Ståhl N, Falkman G, Karlsson A, Mathiason G, Boström J. Deep reinforcement learning for multiparameter optimization in de novo drug design. J Chem Inf Model. 2019;59(7):3166-76. doi:10.1021/acs.jcim.9b00325
15. 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
16. 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

17. 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
18. Fromer JC, Coley CW. Computer-aided multi-objective optimization in small molecule discovery. Patterns (N Y). 2023;4(2):100678. doi:10.1016/j.patter.2023.100678
19. Griffiths RR, Hernández-Lobato JM. Constrained Bayesian optimization for automatic chemical design using variational autoencoders. Chem Sci. 2020;11(3):577-86. doi:10.1039/C9SC04026A
20. Kotsias PC, Arús-Pous J, Chen H, Engkvist O, Tyrchan C, Bjerrum EJ. Direct steering of de novo molecular generation with descriptor conditional recurrent neural networks. Nat Mach Intell. 2020;2(5):254-65. doi:10.1038/s42256-020-0174-5
21. 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
22. Rácz A, Mihalovits LM, Beckers M, Fechner N, Stiefl N, Sirockin F, et al. The changing landscape of medicinal chemistry optimization. Nat Rev Drug Discov. 2025;24(11):870-87. doi:10.1038/s41573-025-01225-1
23. Fang Y, Pan X, Shen HB. De novo drug design by iterative multiobjective deep reinforcement learning with graph-based molecular quality assessment. Bioinformatics. 2023;39(4). doi:10.1093/bioinformatics/btad157
24. Li Y, Zhang L, Liu Z. Multi-objective de novo drug design with conditional graph generative model. J Cheminform. 2018;10(1):33. doi:10.1186/s13321-018-0287-6
25. Blaschke T, Engkvist O, Bajorath J, Chen H. Memory-assisted reinforcement learning for diverse molecular de novo design. J Cheminform. 2020;12(1):68. doi:10.1186/s13321-020-00473-0
26. 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
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. doi:10.3389/fphar.2020.565644
28. Artrith N, Butler KT, Coudert FX, Han S, Isayev O, Jain A, et al. Best practices in machine learning for chemistry. Nat Chem. 2021;13(6):505-8. doi:10.1038/s41557-021-00716-z
29. 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