Pharmacophore

ISSN-2229-5402

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

PERTURBATION ATLASES AS PHARMACOLOGICAL REASONING ENGINES FOR CONNECTING MOLECULAR INTERVENTIONS TO CONTEXT-DEPENDENT CELLULAR RESPONSES

Kadiatou Diallo[1]*, Fode Camara[2], Moriba Keita[3], Sophie Leclercq[4]

1. ***Department of Perturbation Atlases and Pharmacological Reasoning, Faculty of Pharmacy, University of Conakry, Conakry, Guinea.***
2. ***Department of Molecular Interventions and Cellular Responses, Institute of Tropical Medicine, Kindia, Guinea.***
3. ***Department of Context-Dependent Response Modeling, Faculty of Pharmacy, University of N'Zérékoré, N'Zérékoré, Guinea.***
4. ***Department of Cellular Pharmacology and Atlases, Faculty of Pharmaceutical Sciences, University of Montpellier, Montpellier, France.***

ARTICLE INFO

Received:
09 November 2024
Received in revised form:
14 February 2025
Accepted:
19 February 2025
Available online:
28 February 2025

Keywords: Perturbation informatics, Pharmacological reasoning, Cellular response states, Chemical perturbation, Single-cell pharmacology, Mechanism of action

ABSTRACT

Perturbation atlases increasingly connect chemical or genetic interventions with high-dimensional cellular responses, yet their pharmaceutical value remains limited by a persistent reasoning gap. A perturbation signature records what changed under a particular experimental condition, but it does not independently establish why the change occurred, whether it is reproducible across biological contexts, whether it reflects therapeutically achievable exposure, or whether it supports target validation, repurposing, or combination design. This Original Perturbation-Informatics Architecture Article proposes the perturbation reasoning engine as a conceptual architecture for transforming perturbation observations into qualified, context-dependent pharmacological hypotheses. The proposed construct treats each observation as a structured relation among an intervention, dose or perturbation strength, exposure duration, baseline cell state, biological model, assay context, response distribution, and provenance. It then separates representation from subsequent reasoning functions, including directional inference, mechanism attribution, interaction analysis, resistance interpretation, cross-platform harmonization, uncertainty qualification, and bounded decision use. The central contribution is an information architecture in which no single predictive or similarity metric is treated as sufficient evidence of pharmaceutical usefulness. Instead, claims are qualified according to their evidentiary burden, context dependence, transportability, and susceptibility to alternative explanations. The architecture is intended to organize theory building, benchmark design, experimental prioritization, and evidence reporting rather than to constitute a validated predictive system. Its principal limitations include incomplete perturbation coverage, uncertain intracellular exposure, non-equivalence between genetic and pharmacological interventions, latent-state ambiguity, platform shift, and the absence of prospective confirmation for many computationally inferred relationships. Perturbation atlases may therefore become more useful as pharmacological reasoning resources when they preserve context, expose uncertainty, distinguish association from mechanism, and return computational hypotheses to discriminating experiments.

To Cite This Article: Diallo K, Camara F, Keita M, Leclercq S. Perturbation Atlases as Pharmacological Reasoning Engines for Connecting Molecular Interventions to Context-Dependent Cellular Responses. Pharmacophore. 2025;16(1):91-102. https://doi.org/10.51847/A1xH4XUW5O

Introduction

Large-scale perturbation profiling has altered the evidentiary landscape of computational pharmaceutical science by making it possible to compare cellular responses across extensive collections of molecular and genetic interventions. The next-generation Connectivity Map demonstrated how standardized transcriptional profiling could organize a large perturbational space in which compounds, genes, and biological states could be compared through response signatures [1]. This achievement created a powerful data substrate, but it also exposed a conceptual limitation that remains unresolved: a response profile is not itself pharmacological knowledge. Similarity between two signatures may identify shared downstream effects, convergent stress

Corresponding Author: Kadiatou Diallo; Department of Perturbation Atlases and Pharmacological Reasoning, Faculty of Pharmacy, University of Conakry, Conakry, Guinea. E-mail: kadiatou.diallo@ugc.edu.gn.

programs, related targets, or technical artifacts. Conversely, dissimilar signatures may arise from different doses, exposure durations, baseline states, or assay platforms even when interventions act through related mechanisms. The scientific problem is therefore not simply how to predict or retrieve a response, but how to reason from an observed perturbation to a bounded claim about directionality, mechanism, interaction, resistance, or decision relevance.

The usability of perturbation resources also depends on whether observations can be discovered together with the metadata needed to interpret them. The LINCS Data Portal 2.0 illustrates that perturbation signatures require queryable access, standardized descriptors, and provenance-linked metadata before they can be reliably located and compared [2]. Data access, however, should not be mistaken for data suitability. A signature may be technically available while lacking the intervention identity, dose, time, biological model, baseline state, replicate structure, or assay information required for the intended inference. Computational systems that treat all accessible profiles as interchangeable risk converting metadata incompleteness into apparent biological similarity. For pharmaceutical reasoning, the unit of analysis must therefore extend beyond a vector of measured features. It must include the experimental and biological conditions that delimit what the vector can support.

Single-cell perturbation methods have made this requirement more apparent by revealing response heterogeneity that population averages conceal. Massively multiplexed chemical transcriptomics at single-cell resolution showed that chemical interventions can generate distinct responses across cellular subpopulations and states rather than one uniform population-level effect [3]. This heterogeneity matters pharmacologically because treatment sensitivity, adaptive escape, transient stress, lineage-specific vulnerability, and resistant-state enrichment may coexist within the same nominal model. A mean response can consequently appear stable even when therapeutically important subpopulations expand, disappear, or move into alternative states. Yet single-cell resolution alone does not solve the reasoning problem. It increases the dimensionality and granularity of the evidence while also introducing uncertainty in cell-state annotation, distributional alignment, trajectory inference, and cross-platform comparability.

Genome-scale Perturb-seq further demonstrates that intervention effects can be represented as distributed, information-rich cellular programs rather than as isolated endpoint changes [4]. The opportunity created by such atlases is therefore broader than perturbation prediction: they may support a structured form of pharmacological reasoning in which interventions are connected to context-conditioned response states, competing mechanisms, interaction hypotheses, and experimentally testable decisions. This article proposes the perturbation reasoning engine as an original conceptual architecture for that purpose. The contribution is not a validated algorithm, software platform, or operational decision tool. It is an information and inference architecture that separates five scientific tasks that are often conflated: representing the perturbation observation, contextualizing the response, learning directional or mechanistic relations, qualifying cross-platform transport and uncertainty, and restricting application claims to the evidence available. The central argument is that no single measure of predictive performance, response similarity, pathway enrichment, or synergy can establish pharmaceutical usefulness in isolation.

From Perturbation Datasets to Pharmacological Knowledge

A perturbation dataset becomes pharmacologically informative only when the observed response is interpreted relative to the context in which it was generated. High-throughput transcriptional profiling has shown that anti-cancer drugs can elicit both shared and cell-type-specific response programs [5]. This finding means that the same compound cannot be assigned one context-free response identity. Some effects may be conserved across models, whereas others may depend on lineage, genotype, baseline pathway activity, differentiation state, or culture conditions. The relevant knowledge object is therefore not "compound X produces signature Y," but "compound X, at a specified perturbation intensity and duration, produces a particular distribution of responses in a defined biological and technical context." Context should not be treated as an optional covariate added after modeling. It is part of the claim itself because it determines where the response was observed and where it may plausibly be transported.

Signature-search tools demonstrate how perturbation profiles can be transformed into interpretable hypotheses about drug similarity, pathway convergence, or reversal of a disease-associated state. L1000FWD, for example, supports retrieval and visualization of relationships among drug-induced transcriptomic signatures [6]. Such systems are valuable for hypothesis generation because they allow investigators to navigate response spaces that would otherwise remain inaccessible. Nevertheless, signature similarity is mechanistically non-unique. Two interventions may converge on common stress, cell-cycle, apoptotic, or compensatory programs despite acting through different proximal mechanisms. A reversal relation may also arise at an exposure that is not achievable, in a cell model that poorly represents the disease, or through a toxic rather than therapeutically useful response. The proposed architecture therefore treats signature retrieval as an evidence-navigation function, not as a mechanism-confirmation function. It produces candidates for further reasoning, while the strength and meaning of each candidate depend on dose, time, state, orthogonal evidence, and uncertainty.

Representing compounds only by chemical structure is similarly insufficient when the objective is to connect molecular interventions to cellular responses. The Chemical Checker extends small-molecule similarity across chemical, target, network, cellular, and clinical evidence spaces, illustrating that compound relations can be described at multiple biological levels [7]. In the proposed perturbation reasoning engine, these levels are not collapsed into one universal embedding. Instead, they are retained as partially independent evidence channels. Chemical similarity may support a structural analogy; target annotations may support a proximal mechanism hypothesis; network information may identify pathway-level convergence; and cellular signatures may support response-state similarity. Disagreement among these channels is scientifically informative. A structurally dissimilar compound with a similar cellular response may reveal pathway convergence or nonspecific stress,

whereas a structurally similar pair with divergent responses may expose context dependence, differential exposure, or polypharmacology. The engine should preserve these alternatives rather than forcing them into a single score.

Large-scale viability profiling also shows both the promise and the interpretive risk of converting perturbation data into repurposing hypotheses. Systematic profiling of non-oncology drugs has identified unexpected anti-cancer activities while simultaneously highlighting risks related to compound identity, potency, and off-target interpretation [8]. A perturbation atlas may therefore nominate a compound for further investigation, but it cannot independently establish that the compound acts through its annotated target, reaches the required exposure in the intended tissue, possesses an acceptable safety margin, or offers therapeutic benefit. Pharmacological knowledge emerges only when response evidence is connected to traceable intervention identity, biological context, plausible mechanism, exposure feasibility, and discriminating confirmation. Within the proposed architecture, the transition from dataset to knowledge is consequently represented as a graded evidentiary transformation rather than a binary conversion. **Table 1** organizes the evidence, constructs, and interpretation boundaries needed to develop from perturbation datasets to pharmacological knowledge within the article's central argument.

Table 1. Evidence domains, core questions, scientific requirements, and interpretation boundaries for From perturbation datasets to pharmacological knowledge

Construct or evidence domain	Core question	Scientific basis	Role in the article	Failure or overclaiming risk	Boundary statement
Perturbation identity and provenance	What intervention was actually applied, and can its identity and handling history be traced?	Compound identifiers, structures, genetic reagents, batch records, library annotations, perturbation modality, and source metadata	Establishes the intervention as an auditable evidence object rather than an informal label	Misidentified compounds, reagent inconsistency, or annotation errors are interpreted as biological findings	A recorded intervention identity does not establish target engagement or mechanism
Biological response signature	What molecular or phenotypic changes were observed after intervention?	Transcriptomic, proteomic, epigenomic, morphological, viability, or functional measurements	Supplies the empirical response representation that initiates reasoning	A high-dimensional signature is treated as a complete explanation of drug action	A response signature documents association under specified conditions, not causal mechanism
Biological context	In which cell type, cell state, genotype, tissue model, or environmental condition was the response generated?	Cell annotations, lineage, baseline state, disease features, culture system, microenvironment, and donor or model characteristics	Defines the domain within which the response claim is interpretable	Context is treated as a nuisance variable or ignored during transport to another model	A response is not context-independent merely because it is reproducible within one system
Dose, perturbation strength, and time	Under what intensity and duration did the response occur?	Nominal concentration, genetic perturbation efficiency, exposure duration, sampling time, and available exposure proxies	Distinguishes transient, sustained, weak, toxic, and adaptive responses	Compound labels are compared without accounting for exposure conditions	Nominal dose does not automatically represent intracellular exposure or clinically achievable concentration
Cross-level compound representation	Do structural, target, pathway, cellular, and higher-level evidence support the same interpretation?	Chemical descriptors, target annotations, network relations, phenotypic signatures, and clinical or development knowledge	Allows concordance and disagreement across evidence channels to inform reasoning	Multiple evidence levels are collapsed into one opaque similarity score	Agreement strengthens a hypothesis but does not replace experimental discrimination
Response heterogeneity	Does the intervention affect all cells similarly, or does it produce multiple state-specific outcomes?	Single-cell distributions, subpopulation frequencies, rare-state responses, and resistant-state enrichment	Prevents mean responses from concealing therapeutically relevant subpopulations	Population averages are interpreted as uniform cellular effects	Distributional differences require replication and do not independently establish lineage or resistance mechanism
Mechanistic qualification	Which proximal and downstream explanations remain plausible after the response is observed?	Target evidence, genetic perturbations, pathway activity, rescue experiments, orthogonal assays, and temporal ordering	Converts response similarity into ranked, falsifiable mechanism hypotheses	Post-hoc pathway labels are presented as causal explanations	Mechanism remains provisional until competing explanations are experimentally distinguished
Application readiness	What decision can the evidence responsibly support?	Evidence completeness, uncertainty, transportability, exposure plausibility, safety information, and human review	Restricts atlas outputs to appropriate uses such as hypothesis ranking or experiment prioritization	Benchmark performance or signature similarity is equated with therapeutic value	Perturbation evidence alone does not establish clinical utility, regulatory acceptability, or deployment readiness

Representing Compounds, Doses, Times, Cell States, and Response Contexts

The proposed reasoning engine begins with a structured perturbation observation rather than an unqualified response matrix. Multiplexed single-cell transcriptional response profiling has demonstrated that cancer vulnerabilities and mechanism-related programs can differ among individual cells exposed to the same intervention [9]. Accordingly, the minimum representation must distinguish the perturbation object from the response object. The perturbation object contains traceable intervention identity, chemical structure or genetic sequence where applicable, perturbation modality, nominal target evidence, formulation or vehicle information when available, and provenance. The response object contains measured post-intervention states together with uncertainty and replicate structure. Between them lies a context tuple that includes dose or perturbation strength,

exposure duration, sampling time, biological model, baseline cell state, assay platform, and environmental conditions. This separation prevents the system from assigning one permanent response identity to a compound or genetic intervention.

High-throughput chemical perturbation platforms further demonstrate that intervention identity must remain bound to concentration, duration, and biological context throughout analysis. Massively multiplexed chemical profiling with Microwell-seq 2.0 illustrates how large perturbation experiments can combine many compounds and cellular observations, thereby making explicit metadata essential for correct interpretation [10]. In the proposed architecture, dose and time are not merely additional input columns. They define the intervention actually experienced by the system. A low exposure may produce partial target modulation, a higher exposure may recruit secondary targets or generalized stress, and different sampling times may capture immediate signaling, transcriptional adaptation, cell-state transition, or selective survival. Because nominal concentration may differ from intracellular exposure, the architecture should record whether exposure, uptake, free concentration, or target engagement was measured, estimated, or unavailable. Missing exposure information should widen uncertainty rather than be silently treated as evidence of equivalent pharmacological intensity.

Baseline cell state is equally important because a perturbation response is a transformation of an existing biological condition, not an isolated endpoint. scGen demonstrated that latent generative models can learn and predict single-cell perturbation responses across cellular contexts [11]. This supports the usefulness of state-conditioned representation, but it does not imply unrestricted generalization. A latent space may encode biologically meaningful variation, technical variation, or both, and a predicted response may be reliable only within contexts represented sufficiently during training. The proposed engine therefore treats baseline state as a probabilistic, multimodal description that may include transcriptional programs, protein expression, chromatin accessibility, pathway activity, morphology, lineage, genotype, and environmental features. Nominal cell type remains useful, but it is not assumed to be the finest pharmacologically relevant unit. Two cells assigned the same type may occupy different activation, differentiation, stress, or resistance states and consequently respond differently to the same intervention.

Complex perturbation models show why intervention, covariate, and cellular-state representations must remain compositionally explicit. Models developed to predict cellular responses to unseen perturbations and condition combinations rely on separating perturbation identity from contextual covariates rather than treating every observed condition as an unrelated class [12]. The architectural implication is that compounds, doses, times, cell states, and response contexts should be represented as linked but non-interchangeable components. This modularity enables questions such as whether a response is conserved across cell states, whether a dose effect is transported across models, or whether a combination produces a state transition not predictable from either single agent. However, compositional representation does not guarantee correct extrapolation; predictions outside observed perturbation or biological domains require explicit qualification and prospective testing.

Population-level cellular responses should also be modeled as distributions rather than only as average changes. Neural optimal transport has been used to learn single-cell perturbation responses by relating pre- and post-intervention cellular distributions [13]. Such approaches provide a useful basis for representing heterogeneous state transitions, including the appearance, depletion, or redistribution of subpopulations. Yet a mathematically inferred transport plan is not automatically a lineage map, causal trajectory, or mechanistic explanation. Multiple transitions may fit the same observed marginal distributions, particularly when sampling is sparse or intermediate states are absent. The proposed perturbation reasoning engine therefore records three distinct objects: the observed pre- and post-perturbation distributions, the inferred mapping between them, and the assumptions under which that mapping was estimated. This distinction preserves the value of distributional modeling while preventing inferred correspondence from being presented as experimentally observed cellular history.

Architecture of the Perturbation Reasoning Engine

The proposed architecture begins with a probabilistic representation layer that separates biological variation from measurement processes while preserving the conditions under which each observation was generated. Deep generative modeling for single-cell transcriptomics demonstrates how latent-variable models can account for technical noise and biological heterogeneity within a unified probabilistic formulation [14]. In the perturbation reasoning engine, however, the latent representation is not treated as a self-sufficient explanation of cellular response. It functions as an evidence-preserving interface between raw measurements and subsequent reasoning. Each latent response remains linked to the intervention, dose, time, baseline state, biological model, assay platform, replicate structure, and provenance. This linkage is essential because a compact representation can otherwise conceal the specific conditions that make an inferred relation valid or invalid.

A second architectural component integrates multiple response modalities without assuming that they are interchangeable. Joint probabilistic modeling of transcript and protein measurements illustrates how shared biological factors can be learned while retaining modality-specific measurement processes [15]. The proposed engine generalizes this principle to transcriptomic, proteomic, chromatin, morphological, viability, and functional readouts. Multimodal evidence is organized through partially shared latent factors, but discordance is preserved rather than averaged away. A transcriptional response without a corresponding protein or phenotypic change may indicate temporal delay, compensatory regulation, limited assay sensitivity, or an incomplete mechanism hypothesis. Conversely, concordant evidence across modalities may increase confidence without establishing causality. The architecture therefore produces modality-qualified response objects rather than one undifferentiated score.

A third component incorporates biological prior knowledge to support reasoning about perturbations that have not been directly observed in every context. GEARS shows how graph-based information can contribute to prediction of transcriptional outcomes for novel multigene perturbations [16]. Within the proposed architecture, graph priors may encode target relationships, regulatory interactions, pathway membership, genetic dependencies, or compound–protein associations. These priors constrain the search space and can support compositional reasoning, but they are not assumed to be complete or context-neutral. A network derived from one tissue or state may omit interactions that become relevant in another. Consequently, every graph-informed inference retains the provenance, biological domain, and evidence quality of the prior knowledge that produced it. Predicted relations that depend strongly on uncertain graph edges should be flagged for discriminating experiments rather than presented as mechanistic conclusions.

The architecture culminates in a reasoning layer that combines context-specific regulatory inference with reference mapping and explicit uncertainty qualification. CellOracle illustrates how network inference and in silico perturbation can be used to propose context-specific changes in cell identity, while still requiring experimental testing of those simulated transitions [17]. Transfer-learning approaches such as scArches show how new single-cell observations can be mapped to reference atlases while managing technical differences and preserving established biological structure [18]. The proposed engine combines these functions without equating either network simulation or reference proximity with proof of mechanism. Its outputs are ranked, context-bound hypotheses accompanied by assumptions, alternative explanations, and readiness restrictions. **Table 2** organizes the evidence, constructs, and interpretation boundaries needed to develop architecture of the perturbation reasoning engine within the article's central argument. **Figure 1** presents the perturbation-to-mechanism reasoning engine, showing how the article's key components and boundaries are connected within architecture of the perturbation reasoning engine.

Table 2. Components, relationships, uncertainties, and validation needs within Architecture of the perturbation reasoning engine

Component or layer	Inputs	Core function	Expected output	Uncertainty or limitation	Validation requirement
Perturbation identity layer	Compound structure, genetic reagent, intervention modality, batch identity, nominal target annotations, provenance	Creates a traceable representation of the intervention and separates identity from presumed mechanism	Auditable perturbation object with evidence-linked annotations	Misidentification, incomplete target data, polypharmacology, reagent variability	Chemical verification, reagent quality control, orthogonal target-engagement evidence where relevant
Dose–time layer	Nominal concentration, perturbation strength, exposure duration, sampling time, available exposure or engagement measures	Represents the intensity and temporal position of the intervention	Exposure-qualified perturbation specification	Nominal dose may not equal intracellular exposure; sparse time points may miss transient responses	Replicated dose–response and time-course studies with exposure-relevant measurements
Baseline-state and context encoder	Cell type, lineage, genotype, molecular state, tissue model, culture conditions, assay, donor, batch	Defines the biological and technical domain of the response	Context tuple and baseline-state representation	Annotation uncertainty, hidden state variables, incomplete microenvironment	Independent annotation, multimodal confirmation, external-context testing
Probabilistic response layer	Single-cell or population-level molecular and phenotypic measurements	Separates measurement noise from biological variation and represents heterogeneous outcomes	Response distributions with uncertainty and replicate structure	Latent variables may not be biologically identifiable	Replicate consistency, held-out reconstruction, biological-conservation assessment
Multimodal integration layer	Transcriptomic, proteomic, chromatin, morphological, viability, and functional readouts	Identifies shared and modality-specific response factors	Modality-qualified response representation	Missing modalities, assay-specific noise, incompatible feature scales	Matched multimodal experiments, modality-ablation analysis, concordance testing
Directional inference layer	Pre- and post-perturbation states, time-resolved observations, controls, transport or trajectory models	Estimates the direction and structure of state change under explicit assumptions	Candidate state transitions and directional hypotheses	Sparse sampling, branching trajectories, inferred rather than observed correspondence	Longitudinal sampling, lineage information, perturb-and-rescue experiments
Mechanism reasoning layer	Drug–target evidence, genetic screens, pathway networks, regulatory models, orthogonal assays	Ranks competing proximal and downstream explanations	Mechanism hypotheses with supporting and contradictory evidence	Polypharmacology, network incompleteness, genetic–pharmacological nonequivalence	Target engagement, rescue, counter-perturbation, biochemical and phenotypic confirmation

Synergy and resistance layer	Single-agent profiles, combination dose matrices, longitudinal responses, survivor states	Distinguishes non-additive interactions, adaptive responses, and context-specific escape	Combination and resistance hypotheses	Reference-model dependence, sparse combination space, survivor bias	Replicated dose matrices, temporal validation, mechanistic perturbation, safety and exposure assessment
Harmonization and provenance layer	Platform labels, batch metadata, reference atlases, transformation history	Supports qualified transport across datasets while retaining data lineage	Harmonized evidence objects with traceable transformations	Overcorrection, undercorrection, batch–biology confounding	Matched controls, leave-platform-out tests, biological-conservation metrics
Uncertainty and readiness layer	Calibration data, external contexts, sensitivity analyses, failure audits, intended decision	Relates evidence quality and extrapolation risk to permissible use	Calibrated uncertainty, abstention conditions, readiness classification	Unknown distribution shift and subgroup miscalibration	Prospective calibration, stress testing, human review, use-case-specific decision studies

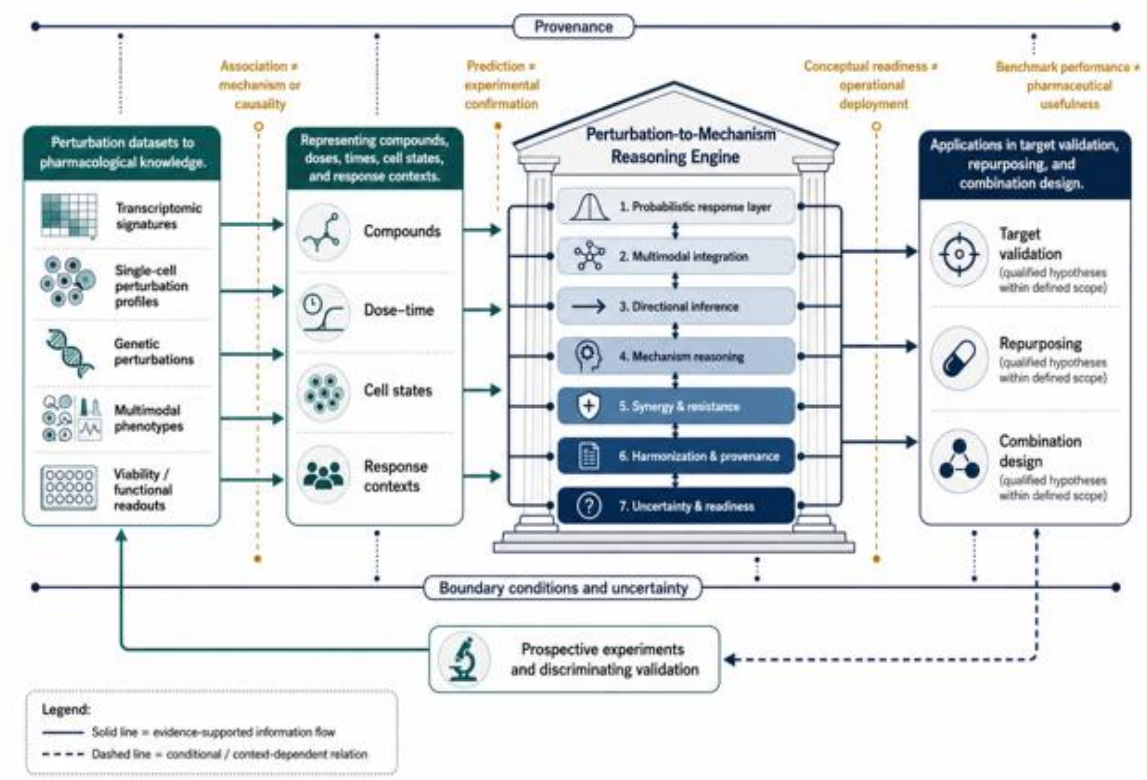

Figure 1. Perturbation-to-Mechanism Reasoning Engine.

The figure is an original conceptual synthesis that organizes the article's central contribution across perturbation datasets to pharmacological knowledge, representing compounds, doses, times, cell states, and response contexts, representing compounds, cell states, response contexts, architecture of the perturbation reasoning engine. 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.

Learning Directionality, Mechanism, Synergy, and Resistance

Directionality, mechanism, synergy, and resistance represent distinct inference problems and should not be collapsed into a general notion of model performance. Rich single-cell phenotypes have been used to construct genetic interaction manifolds in which combinatorial perturbations reveal response programs that cannot be recovered from scalar viability measurements alone [19]. Such data can distinguish whether two perturbations converge, oppose one another, buffer a phenotype, or create a qualitatively new state. Nevertheless, interaction structure does not by itself establish pharmacological synergy. Genetic perturbations may differ from drug action in magnitude, timing, reversibility, selectivity, and subcellular distribution. The reasoning engine therefore separates descriptive interaction geometry from claims about therapeutic combinations.

Learning directionality requires evidence that distinguishes state movement from static resemblance. Optimal-transport analysis of time-resolved single-cell data demonstrates how distributions sampled across a biological process can be linked to infer likely trajectories [20]. In perturbation informatics, analogous models may estimate whether an intervention moves cells toward, away from, or between defined states. Such inference is stronger when observations are temporally ordered,

intermediate states are sampled, and controls constrain alternative transitions. It remains weaker when pre- and post-treatment populations are unpaired or when branching and selective survival can produce similar distributions. Directional inference should therefore be reported together with its temporal assumptions and should not be described as causal merely because an arrow can be estimated.

Mechanism learning requires convergence among evidence types that differ in their proximity to the causal process. Integration of pharmacological and CRISPR screens has shown how genetic dependencies can help interpret drug mechanism of action and distinguish plausible targets from broader response associations [21]. Yet pharmacological and genetic perturbations are not equivalent. A drug may engage several proteins, modulate rather than remove function, act in a compartment-specific manner, or produce exposure-dependent secondary effects. Genetic concordance can strengthen a mechanism hypothesis, but discordance may reflect polypharmacology, incomplete perturbation, compensatory biology, or incorrect target annotation. The proposed engine therefore represents mechanism as a ranked set of alternatives supported by target, genetic, temporal, pathway, and phenotypic evidence.

Combination and resistance reasoning require evaluation designs that preserve dose, context, and alternative additivity assumptions. Community assessment of cancer drug-combination prediction demonstrated that performance depends on the dataset, held-out setting, and scoring formulation used to define the prediction task [22]. Large-scale testing across breast, colon, and pancreatic cancer models further showed that effective combinations depend on molecular context and on the relationship among potency, efficacy, and interaction behavior [23]. A predicted combination should consequently remain a hypothesis until tested across an informative dose matrix and evaluated for exposure feasibility, mechanism, and toxicity. Resistance claims similarly require evidence that a state precedes or causally contributes to reduced response, rather than merely appearing among surviving cells. **Table 3** organizes the evidence, constructs, and interpretation boundaries needed to develop learning directionality, mechanism, synergy, and resistance within the article's central argument.

Table 3. Components, relationships, uncertainties, and validation needs within Learning directionality, mechanism, synergy, and resistance

Evaluation dimension	What must be tested	Suitable evidence or assessment	Failure signal	Interpretive limitation	Decision relevance
Directional state change	Whether the intervention moves cells between defined states in a reproducible direction	Time-resolved perturbations, controls, longitudinal sampling, lineage information, transport or trajectory consistency	Direction reverses across replicates or disappears when intermediate states are included	Direction does not establish causal mechanism and may reflect selection rather than transition	Supports prioritization of temporal and perturb-and-rescue experiments
Temporal coherence	Whether early, intermediate, and sustained responses follow a plausible ordering	Dense time courses, pathway activation sequences, repeated measurements, target-engagement timing	Predicted downstream events precede required upstream events without supporting evidence	Sparse sampling may make several temporal explanations equally plausible	Helps distinguish immediate target effects from adaptation or toxicity
Mechanism attribution	Whether a proposed target or pathway explanation discriminates among alternatives	Genetic perturbation, target engagement, biochemical assays, rescue, counter-perturbation, pathway-specific phenotypes	Mechanism label remains unchanged despite contradictory orthogonal evidence	Concordance may arise through shared downstream programs	Supports selection of experiments designed to falsify competing mechanisms
Genetic–pharmacological concordance	Whether genetic and compound perturbations produce sufficiently aligned context-specific effects	CRISPR or RNA-interference screens, compound profiling, target annotations, dose dependence	Genetic dependency and drug response disagree systematically across contexts	Differences may reflect perturbation strength, reversibility, polypharmacology, or exposure	Qualifies target-validation and mechanism hypotheses
Synergy robustness	Whether a combination remains non-additive across relevant doses, models, and reference assumptions	Replicated dose matrices, multiple additivity models, held-out contexts, efficacy and potency analysis	Interaction changes sign with minor scoring or dose-window changes	In vitro synergy does not establish therapeutic benefit or acceptable safety	Supports preclinical combination prioritization
Mechanistic complementarity	Whether combined interventions affect distinct but interacting vulnerabilities	Pathway evidence, single-agent and combination signatures, rescue studies, target overlap analysis	Combination response is explained entirely by nonspecific toxicity or one dominant agent	Complementary signatures do not guarantee compatible exposure or tolerability	Helps choose combinations for discriminating experimental tests
Adaptive resistance	Whether a post-treatment state reflects an induced escape program	Longitudinal single-cell data, state-frequency change, perturb-and-rescue testing, repeated exposure	Putative resistance state is present before treatment at the same frequency or is not	Association among surviving cells may reflect selection rather than induced adaptation	Supports hypotheses for sequential or combination intervention

			functionally linked to reduced response		
Pre-existing resistance	Whether a baseline subpopulation predicts and contributes to poor response	Pretreatment state profiling, lineage tracing, prospective isolation, functional perturbation	Baseline signature fails to predict response in independent contexts	Predictive enrichment does not prove causal dependence	Supports context-stratified experimental design
Extrapolation to unseen combinations	Whether models generalize to new pairs and biological contexts	Pair-held-out, drug-held-out, and context-held-out validation; uncertainty calibration	Strong random-split performance collapses under drug or context holdout	Benchmark success remains local to the sampled chemical and biological space	Determines whether predictions are suitable for exploratory prioritization
Pharmaceutical relevance	Whether the inferred interaction is compatible with exposure, selectivity, and safety constraints	Pharmacological characterization, achievable exposure, target engagement, toxicity and developability evidence	Predicted interaction requires infeasible exposure or unacceptable nonspecific effects	Perturbation atlases rarely contain sufficient translational evidence alone	Defines the boundary between computational nomination and therapeutic development

Cross-Platform Harmonization and Uncertainty-Aware Inference

Cross-platform harmonization is necessary because perturbation atlases are generated using different assays, laboratories, feature spaces, normalization procedures, and biological models. Comparative assessment of batch-effect correction methods in single-cell RNA sequencing has shown that integration methods vary in their ability to remove technical variation while preserving biological differences [24]. In a perturbation reasoning engine, harmonization should therefore be evaluated against two competing errors: undercorrection, which leaves technical artifacts that resemble context-specific biology, and overcorrection, which erases real response differences. Visual mixing of batches is insufficient because a fully mixed representation may have removed biologically meaningful perturbation structure. Harmonization must be tied to the intended comparison and supported by controls that distinguish platform effects from intervention effects.

Atlas-level integration introduces further challenges because the complexity of the task changes with the number of batches, cell states, modalities, and biological differences involved. Benchmarking of atlas-level integration methods demonstrates that no single metric adequately captures batch removal, conservation of biological structure, scalability, and behavior across task types [25]. The proposed architecture therefore uses a multidimensional integration audit rather than one harmonization score. Each transformed dataset should retain its original identifiers, transformation history, and diagnostics. Performance should be examined separately for abundant and rare states, controls and perturbed samples, familiar and unseen contexts, and local and global structure. An integrated representation is considered usable only for the specific inference tasks for which biological conservation and technical correction have been demonstrated.

Multimodal integration should similarly preserve evidence that is unique to each measurement process. MOFA+ provides a statistical framework for learning shared and modality-specific factors from multimodal single-cell data [26]. This principle is important because missing or discordant modalities can contain information about mechanism and uncertainty. A transcriptomic factor may align with protein abundance in one context but not another; a morphological phenotype may persist after an early transcriptional response has resolved. The reasoning engine therefore avoids treating multimodal integration as compulsory feature fusion. Instead, it records which conclusions depend on shared factors, which depend on one modality, and which are weakened by missing data. Latent factors may organize evidence, but they should not be assigned causal meaning without additional validation.

Uncertainty-aware inference is the final requirement for responsible cross-context transport. Reviews of uncertainty quantification in deep learning distinguish uncertainty associated with limited knowledge from variability inherent in the data and emphasize that both require explicit treatment [27]. A broader survey of uncertainty in deep neural networks further shows that uncertainty estimates can remain miscalibrated, particularly under distribution shift [28]. The proposed engine therefore separates data uncertainty, model uncertainty, annotation uncertainty, harmonization uncertainty, mechanism ambiguity, and decision uncertainty. It also permits abstention when an intervention, cell state, platform, or combination lies outside the supported domain. Model confidence is not treated as calibrated uncertainty, and calibrated uncertainty is not treated as evidence of mechanism or usefulness. Instead, uncertainty constrains the strength of the claim and the next experimental action.

Applications in Target Validation, Repurposing, and Combination Design

In target validation, perturbation atlases may help prioritize genes or pathways whose disruption produces a context-specific phenotype of interest. Genome-scale CRISPR screens have been used to prioritize cancer therapeutic targets by integrating dependency effects with genomic context and tractability considerations [29]. The reasoning engine extends this logic by requiring that each target hypothesis specify the responsive cell state, relevant model, perturbation modality, effect direction, and uncertainty. A genetic dependency may justify orthogonal validation, but it does not establish that a target can be modulated selectively with a drug, that the required exposure is achievable, or that inhibition will be tolerated in normal tissues. Target validation is therefore represented as a staged evidentiary process rather than a binary label.

Drug repurposing provides a second application because perturbation signatures can connect existing compounds to disease-associated states or mechanisms. Reviews of repurposing practice emphasize that promising hypotheses must be evaluated

together with pharmacology, safety, exposure, intellectual-property, formulation, and clinical-development considerations [30]. Within the proposed architecture, signature reversal or state normalization is only the beginning of the reasoning chain. The system must determine whether the compound identity is reliable, whether the relevant effect occurs at a plausible exposure, whether the inferred target is supported, and whether the disease model captures the intended context. Repurposing readiness therefore depends on evidence not usually contained in a perturbation atlas, and computational nomination remains a research hypothesis rather than a therapeutic recommendation.

Resource quality also affects whether repurposing hypotheses can be experimentally tested. The Drug Repurposing Hub illustrates the value of curated compound identity, target annotation, clinical status, and accessible physical libraries for systematic screening [31]. These attributes allow perturbation observations to be connected to traceable interventions and follow-up experiments. The reasoning engine treats such annotation as a provenance layer rather than as proof that the listed mechanism is correct. A compound may have several targets, and its observed response may arise through a different target at the tested concentration. Repurposing hypotheses should consequently retain alternative mechanisms and require orthogonal testing, including target engagement, dose–response characterization, and comparison with structurally or mechanistically informative controls.

Combination design is the most demanding application because it expands the intervention space while adding dependence on dose, timing, reference model, and resistance context. SynergyFinder 2.0 illustrates how combination sensitivity and synergy can be calculated and benchmarked using alternative interaction formulations [32]. These formulations are useful analytical tools, but their scores are conditional on the selected dose matrix and null model. Cancer dependency maps likewise support context-specific hypotheses about vulnerabilities while requiring subsequent pharmacological validation and selectivity assessment [33]. The proposed reasoning engine therefore connects combination hypotheses to the underlying single-agent responses, candidate mechanisms, context-specific dependencies, uncertainty, and feasibility constraints. **Table 4** organizes the evidence, constructs, and interpretation boundaries needed to develop applications in target validation, repurposing, and combination design within the article's central argument.

Table 4. Evaluation, implementation, and research priorities arising from Perturbation Atlases as Pharmacological Reasoning Engines for Connecting Molecular Interventions to Context-Dependent

Research or implementation priority	Unresolved gap	Required methodological work	Evidence needed	Responsible actors	Expected contribution
Minimum perturbation ontology	Compound, dose, time, state, assay, and provenance are inconsistently represented	Develop machine-readable minimum fields, controlled vocabularies, and cross-resource mappings	Metadata audits, expert consensus, inter-resource mapping studies	Data curators, pharmacologists, bioinformaticians, ontology developers	Interoperable and auditable perturbation evidence objects
Exposure-aware response modeling	Nominal concentration is often disconnected from intracellular exposure and target engagement	Link response profiles to free concentration, uptake, engagement, and temporal exposure measures where possible	Matched exposure, target-engagement, and molecular-response experiments	Experimental pharmacologists, pharmacometricians, assay scientists	More pharmacologically coherent dose–response interpretation
Longitudinal state-transition atlases	Most atlases contain sparse snapshots rather than observed transitions	Generate dense time courses with lineage-aware or repeated state measurements	Longitudinal single-cell and multimodal perturbation studies	Cell biologists, systems pharmacologists, technology developers	Stronger distinction among transient response, adaptation, selection, and resistance
Mechanism-discriminating benchmarks	Prediction tasks rarely test whether competing mechanisms can be separated	Design benchmarks containing alternative hypotheses, negative controls, rescues, and counter-perturbations	Prospective experiments capable of falsifying specific mechanisms	Benchmark consortia, experimentalists, computational scientists	Evaluation that separates predictive fit from mechanistic discrimination
Cross-platform reference standards	Harmonization lacks matched materials and perturbations across technologies	Establish shared reference samples, controls, intervention panels, and provenance requirements	Inter-laboratory and inter-platform ring studies	Standards organizations, consortia, laboratories, publishers	Direct assessment of technical correction and biological conservation
Context-stratified uncertainty	Average calibration can hide high-confidence errors in rare or shifted states	Develop uncertainty and abstention methods stratified by intervention novelty, cell state, and platform	External validation under pre-specified distribution shifts	Machine-learning researchers, statisticians, domain scientists	Uncertainty estimates that better reflect extrapolation and decision risk
Target-validation workflow	Genetic dependency, mechanism, tractability, and pharmacological feasibility are often evaluated separately	Link perturbation evidence to staged orthogonal validation and target-product-profile questions	Genetic, pharmacological, biochemical, safety, and normal-tissue evidence	Target biologists, medicinal chemists, toxicologists, project teams	More transparent prioritization without equating dependency with druggability
Repurposing qualification	Signature matching is frequently disconnected from exposure, safety, and development constraints	Integrate perturbation evidence with compound quality, clinical status, pharmacokinetics,	Curated compound records and prospective disease-relevant testing	Pharmacologists, clinicians, formulation scientists, translational teams	Repurposing hypotheses with explicit feasibility and evidence boundaries

		formulation, and safety information			
Combination and resistance evaluation	Synergy scores are sensitive to dose design, reference model, and biological context	Use replicated dose matrices, longitudinal profiling, alternative null models, and resistance-focused experiments	Single-agent and combination response, exposure, toxicity, and state-transition data	Combination pharmacologists, systems biologists, toxicologists	More robust prioritization of combinations and resistance-informed follow-up
Human oversight and governance	Use roles, override conditions, documentation, and escalation criteria are underdefined	Establish use-case-specific review procedures, audit trails, and stop conditions	Human-factors studies, decision audits, failure simulations	Research leaders, quality teams, domain experts, governance bodies	Accountable use without implying autonomous or deployment-ready decision making

Limitations and Boundary Conditions

The proposed perturbation reasoning engine is a conceptual architecture and has not been empirically validated as an integrated system. Its modules are supported by existing perturbation resources, predictive models, integration methods, and application studies, but their combination into one reasoning architecture remains a proposed synthesis. The architecture does not establish that all required metadata can be collected consistently or that heterogeneous perturbation platforms can be harmonized without loss. It also does not demonstrate superiority over existing analytic workflows. Future evaluation must therefore examine whether the proposed separation of representation, inference, uncertainty, and decision use produces more reliable scientific conclusions than simpler approaches.

The available evidence base is also uneven across intervention types, biological contexts, and response modalities. Cell lines and simplified culture systems may not reproduce tissue architecture, immune interactions, metabolism, distribution, or chronic exposure. Genetic perturbations may not reproduce the magnitude, reversibility, selectivity, or pharmacokinetics of molecular interventions. Nominal compound concentration may not represent free intracellular exposure or target engagement. Single-cell measurements may introduce dissociation effects, sparse sampling, annotation uncertainty, and platform-specific biases. Consequently, transport from one perturbation atlas to another, or from experimental models to therapeutic contexts, remains conditional and should widen uncertainty rather than be assumed.

The architecture cannot independently establish causality, mechanism, therapeutic benefit, safety, developability, clinical utility, regulatory acceptability, or deployment readiness. It cannot convert an association into a mechanism merely by adding a network explanation, nor can it convert benchmark performance into pharmaceutical usefulness. It may help organize evidence and prioritize discriminating experiments, but it should not replace biochemical confirmation, exposure characterization, functional validation, safety assessment, or expert judgment. Misuse is most likely when a single response score, similarity measure, pathway attribution, uncertainty value, or synergy estimate is detached from its biological and decision context.

Research Agenda and Implementation Priorities

The first research priority is to establish whether a minimum perturbation representation improves reproducibility and interpretability across resources. This requires prospective studies in which intervention identity, dose, time, baseline state, assay, response distribution, and provenance are recorded using shared definitions. Evaluation should test whether independent groups can reconstruct the same evidence object and reach similar conclusions about where a response is transportable. Reference materials and matched perturbations should be profiled across laboratories and platforms to distinguish technical correction from biological conservation. Reporting should include missingness, transformation history, excluded contexts, and the assumptions used for integration.

The second priority is to develop benchmarks that separate prediction from directionality, mechanism, synergy, and resistance. Random train–test splits and aggregate prediction scores are insufficient for this purpose. Evaluation designs should include intervention-held-out, cell-state-held-out, platform-held-out, and laboratory-held-out settings, together with prospective perturbations that can discriminate among competing explanations. Mechanism benchmarks should include rescue and counter-perturbation experiments. Combination benchmarks should report sensitivity to dose range and additivity model. Resistance studies should distinguish pre-existing states from induced adaptation through longitudinal or lineage-aware designs. Uncertainty should be evaluated by whether it predicts actual error and supports appropriate abstention under distribution shift.

Implementation should proceed in stages rather than through immediate deployment. An initial stage should focus on data representation, provenance, and retrospective evidence audits. A second stage should evaluate individual modules in bounded research tasks, such as selecting follow-up perturbations or identifying contexts in which a mechanism hypothesis is most uncertain. A third stage should test the integrated architecture prospectively with pre-specified hypotheses, alternatives, decision thresholds, and human-review procedures. Progression between stages should depend on demonstrated calibration, reproducibility, biological conservation, and scientific utility. Operational use should remain contingent on the intended decision, available oversight, and evidence that the system improves reasoning without concealing uncertainty or displacing necessary experimental confirmation.

Conclusion

Perturbation atlases can contribute more to pharmaceutical science when they are treated not merely as collections of response signatures but as structured evidence systems that connect interventions to context-dependent cellular changes. The proposed perturbation reasoning engine provides an original conceptual architecture for organizing that transition. It separates intervention identity, dose, time, baseline state, response context, multimodal representation, directional inference, mechanism attribution, synergy, resistance, harmonization, uncertainty, and decision use. Its central premise is that no single similarity, prediction, pathway, or interaction metric can establish pharmaceutical usefulness. Each output must remain linked to its evidentiary basis, biological domain, assumptions, alternatives, and readiness boundary. The architecture is intended to support theory building, reporting, benchmark design, and experimental prioritization, not to claim empirical validation or autonomous decision authority. Its scientific value will depend on prospective falsification, context-stratified evaluation, exposure-aware modeling, cross-platform reference standards, and human oversight.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Subramanian A, Narayan R, Corsello SM, Peck DD, Natoli TE, Lu X, et al. A next generation connectivity map: L1000 platform and the first 1,000,000 profiles. Cell. 2017;171(6):1437-52. doi:10.1016/j.cell.2017.10.049
2. Stathias V, Turner J, Koleti A, Vidovic D, Cooper DJ, Fazel-Najafabadi M, et al. LINCS Data Portal 2.0: Next generation access point for perturbation-response signatures. Nucleic Acids Res. 2020;48(D1). doi:10.1093/nar/gkz1023
3. Srivatsan SR, McFaline-Figueroa JL, Ramani V, Saunders L, Cao J, Packer JS, et al. Massively multiplex chemical transcriptomics at single-cell resolution. Science. 2020;367(6473):45-51. doi:10.1126/science.aax6234
4. Replogle JM, Saunders RA, Pogson AN, Hussmann JA, Lenail A, Guna A, et al. Mapping information-rich genotype-phenotype landscapes with genome-scale Perturb-seq. Cell. 2022;185(14):2559-75. doi:10.1016/j.cell.2022.05.013
5. Niepel M, Hafner M, Duan Q, Wang Z, Paull EO, Chung M, et al. Common and cell-type specific responses to anti-cancer drugs revealed by high throughput transcript profiling. Nat Commun. 2017;8(1):1186. doi:10.1038/s41467-017-01383-w
6. Wang Z, Lachmann A, Keenan AB, Ma'ayan A. L1000FWD: Fireworks visualization of drug-induced transcriptomic signatures. Bioinformatics. 2018;34(12):2150-2. doi:10.1093/bioinformatics/bty060
7. Duran-Frigola M, Pauls E, Guitart-Pla O, Bertoni M, Alcalde V, Amat D, et al. Extending the small-molecule similarity principle to all levels of biology with the Chemical Checker. Nat Biotechnol. 2020;38(9):1087-96. doi:10.1038/s41587-020-0502-7
8. Corsello SM, Nagari RT, Spangler RD, Rossen J, Kocak M, Bryan JG, et al. Discovering the anti-cancer potential of non-oncology drugs by systematic viability profiling. Nat Cancer. 2020;1(2):235-48. doi:10.1038/s43018-019-0018-6
9. McFarland JM, Paolella BR, Warren A, Geiger-Schuller K, Shibue T, Rothberg M, et al. Multiplexed single-cell transcriptional response profiling to define cancer vulnerabilities and therapeutic mechanism of action. Nat Commun. 2020;11(1):4296. doi:10.1038/s41467-020-17440-w
10. Chen H, Liao Y, Zhang G, Sun Z, Yang L, Fang X, et al. High-throughput Microwell-seq 2.0 profiles massively multiplexed chemical perturbation. Cell Discov. 2021;7(1):107. doi:10.1038/s41421-021-00333-7
11. Lotfollahi M, Wolf FA, Theis FJ. scGen predicts single-cell perturbation responses. Nat Methods. 2019;16(8):715-21. doi:10.1038/s41592-019-0494-8
12. Lotfollahi M, Klimovskaia Susmelj A, De Donno C, Hetzel L, Ji Y, Ibarra IL, et al. Predicting cellular responses to complex perturbations in high-throughput screens. Mol Syst Biol. 2023;19(6). doi:10.15252/msb.202211517
13. Bunne C, Stark SG, Gut G, Sarabia Del Castillo J, Levesque M, Lehmann KV, et al. Learning single-cell perturbation responses using neural optimal transport. Nat Methods. 2023;20(11):1759-68. doi:10.1038/s41592-023-01969-x
14. Lopez R, Regier J, Cole MB, Jordan MI, Yosef N. Deep generative modeling for single-cell transcriptomics. Nat Methods. 2018;15(12):1053-8. doi:10.1038/s41592-018-0229-2
15. Gayoso A, Steier Z, Lopez R, Regier J, Nazor KL, Streets A, et al. Joint probabilistic modeling of single-cell multi-omic data with totalVI. Nat Methods. 2021;18(3):272-82. doi:10.1038/s41592-020-01050-x
16. Roohani Y, Huang K, Leskovec J. Predicting transcriptional outcomes of novel multigene perturbations with GEARS. Nat Biotechnol. 2024;42(6):927-35. doi:10.1038/s41587-023-01905-6
17. Kamimoto K, Stringa B, Hoffmann CM, Jindal K, Solnica-Krezel L, Morris SA. Dissecting cell identity via network inference and in silico gene perturbation. Nature. 2023;614(7949):742-51. doi:10.1038/s41586-022-05688-9
18. Lotfollahi M, Naghipourfar M, Luecken MD, Khajavi M, Büttner M, Wagenstetter M, et al. Mapping single-cell data to reference atlases by transfer learning. Nat Biotechnol. 2022;40(1):121-30. doi:10.1038/s41587-021-01001-7

19. Norman TM, Horlbeck MA, Replogle JM, Ge AY, Xu A, Jost M, et al. Exploring genetic interaction manifolds constructed from rich single-cell phenotypes. Science. 2019;365(6455):786-93. doi:10.1126/science.aax4438
20. Schiebinger G, Shu J, Tabaka M, Cleary B, Subramanian V, Solomon A, et al. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. Cell. 2019;176(4):928-43. doi:10.1016/j.cell.2019.01.006
21. Gonçalves E, Segura-Cabrera A, Pacini C, Picco G, Behan FM, Jaaks P, et al. Drug mechanism-of-action discovery through the integration of pharmacological and CRISPR screens. Mol Syst Biol. 2020;16(7). doi:10.15252/msb.20199405
22. Menden MP, Wang D, Mason MJ, Szalai B, Bulusu KC, Guan Y, et al. Community assessment to advance computational prediction of cancer drug combinations in a pharmacogenomic screen. Nat Commun. 2019;10(1):2674. doi:10.1038/s41467-019-09799-2
23. Jaaks P, Coker EA, Vis DJ, Edwards O, Carpenter EF, Leto SM, et al. Effective drug combinations in breast, colon and pancreatic cancer cells. Nature. 2022;603(7899):166-73. doi:10.1038/s41586-022-04437-2
24. Tran HTN, Ang KS, Chevrier M, Zhang X, Lee NYS, Goh M, et al. A benchmark of batch-effect correction methods for single-cell RNA sequencing data. Genome Biol. 2020;21(1):12. doi:10.1186/s13059-019-1850-9
25. Luecken MD, Büttner M, Chaichoompu K, Danese A, Interlandi M, Mueller MF, et al. Benchmarking atlas-level data integration in single-cell genomics. Nat Methods. 2022;19(1):41-50. doi:10.1038/s41592-021-01336-8
26. Argelaguet R, Arnol D, Bredikhin D, Deloro Y, Velten B, Marioni JC, et al. MOFA+: A statistical framework for comprehensive integration of multi-modal single-cell data. Genome Biol. 2020;21(1):111. doi:10.1186/s13059-020-02015-1
27. Abdar M, Pourpanah F, Hussain S, Rezazadegan D, Liu L, Ghavamzadeh M, et al. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. Inf Fusion. 2021;76:243-97. doi:10.1016/j.inffus.2021.05.008
28. Gawlikowski J, Tassi CRN, Ali M, Lee J, Humt M, Feng J, et al. A survey of uncertainty in deep neural networks. Artif Intell Rev. 2023;56(Suppl 1):1513-89. doi:10.1007/s10462-023-10562-9
29. Behan FM, Iorio F, Picco G, Gonçalves E, Beaver CM, Migliardi G, et al. Prioritization of cancer therapeutic targets using CRISPR-Cas9 screens. Nature. 2019;568(7753):511-6. doi:10.1038/s41586-019-1103-9
30. Pushpakom S, Iorio F, Eyers PA, Escott KJ, Hopper S, Wells A, et al. Drug repurposing: Progress, challenges and recommendations. Nat Rev Drug Discov. 2019;18(1):41-58. doi:10.1038/nrd.2018.168
31. Corsello SM, Bittker JA, Liu Z, Gould J, McCarren P, Hirschman JE, et al. The Drug Repurposing Hub: A next-generation drug library and information resource. Nat Med. 2017;23(4):405-8. doi:10.1038/nm.4306
32. Ianevski A, Giri AK, Aittokallio T. SynergyFinder 2.0: Visual analytics of multi-drug combination synergies. Nucleic Acids Res. 2020;48(W1). doi:10.1093/nar/gkaa216
33. Tsherniak A, Vazquez F, Montgomery PG, Weir BA, Kryukov G, Cowley GS, et al. Defining a cancer dependency map. Cell. 2017;170(3):564-76. doi:10.1016/j.cell.2017.06.010