Pharmacophore

ISSN-2229-5402

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

HIERARCHICAL NEURAL NETWORKS FOR DRUG RESPONSE PREDICTION USING TARGETS, TRANSCRIPTOMICS, AND PHARMACOGENOMIC BIOMARKERS

Fatima Al-Zahra[1]*, Amina El Idrissi[1]

1. *Department of AI Pharmaceutical Systems, Faculty of Medicine and Pharmacy, University of Casablanca, Casablanca, Morocco.*

ARTICLE INFO

Received:
15 January 2026
Received in revised form:
26 March 2026
Accepted:
01 April 2026
Available online:
28 April 2026

Keywords: Hierarchical neural network, Drug response prediction, Precision oncology, Pharmacogenomics, Transcriptomics, Multimodal deep learning

ABSTRACT

Precision oncology seeks to tailor anticancer therapies to individual patients by leveraging molecular profiles of tumors and, when available, patient-specific data; however, predictive models trained on large pharmacogenomic screens often struggle to translate clear in vitro associations into clinically meaningful response predictions. Many existing drug response models combine drug descriptors, gene expression levels, mutations, and copy-number alterations into a single flat representation, disregarding the biological sequence in which a drug engages molecular targets, perturbs cellular programs, and is influenced by the patient's genomic context. To address this, this MDL article proposes a hierarchical neural network for predicting drug response in cancer, designed to integrate drug-target interactions, tumor transcriptomes, and pharmacogenomic biomarkers within a biologically ordered framework. The architecture first encodes drug-target and target-inhibition information in a lower-level target-engagement module, then merges this latent drug mechanism representation with transcriptomic context in a middle module, and finally modulates the resulting signal using mutation and copy-number biomarkers in an upper pharmacogenomic module to produce a personalized sensitivity estimate. Conceptually, this nested structure is expected to yield more biologically coherent predictions than flat multimodal approaches and could clarify whether predicted resistance arises primarily from inadequate target engagement, an unfavorable transcriptomic state, or genomic alterations affecting drug response. By respecting the hierarchy from drug mechanism to cellular state to patient genotype, such a biologically structured deep learning model offers a principled framework for interpretable precision oncology, potentially bridging preclinical pharmacogenomic screening and clinical decision support.

To Cite This Article: Al-Zahra F, El Idrissi A. Hierarchical Neural Networks for Drug Response Prediction Using Targets, Transcriptomics, and Pharmacogenomic Biomarkers. Pharmacophore. 2026;17(2):63-72. https://doi.org/10.51847/BP9qbtx7Wl

Introduction

Precision oncology has been shaped by large pharmacogenomic resources that connect molecular profiles with drug sensitivity readouts, enabling deep learning models to predict response from cancer cell line and tumor data [1]. Yet many models still treat drugs, transcriptomes, mutations, and copy-number alterations as parallel inputs that can be concatenated before prediction, even though these variables occupy different biological levels [2]. Integrated genomic-profile models have shown that combining multiple molecular layers can support response prediction, but the biological hierarchy among those layers is often implicit rather than structurally encoded [3]. A model that preserves the order from drug mechanism to cellular state to patient genotype could therefore better match the structure of the biological problem.

Traditional deep learning approaches can learn nonlinear relationships between gene expression and drug sensitivity, but they may obscure the causal sequence by which anticancer drugs act on molecular targets before altering downstream cellular programs [4]. Deep convolutional and fully connected models trained on phenotypic drug response can represent complex associations, yet they often compress biologically distinct mechanisms into a single latent space [5]. Graph and hybrid neural methods such as DeepCDR illustrate the value of representing drugs and cellular features with structured encoders, but the connection between target engagement, transcriptomic context, and pharmacogenomic modulation remains a design challenge [6]. This motivates architectures in which each biological layer has a dedicated module rather than being flattened into a generic feature vector.

Corresponding Author: Fatima Al-Zahra; Department of AI Pharmaceutical Systems, Faculty of Medicine and Pharmacy, University of Casablanca, Casablanca, Morocco. E-mail: fatima.zahra@gmail.com.

Hierarchical neural networks are well suited to domains in which lower-level signals are progressively contextualized by higher-level information. In drug response modeling, DrugCell demonstrated that a neural architecture can be organized around biological knowledge to connect genotype, cellular subsystems, and drug sensitivity in an interpretable manner [7]. Few-shot learning frameworks further show that drug response models can be designed to transfer from large screens to patient settings when the architecture is aligned with the structure of the task [8]. These examples suggest that hierarchical design is not merely a representational preference, but a way to encode assumptions about how drug mechanism, cell phenotype, and patient-specific variation interact.

The central thesis of this article is that drug response prediction should be modeled as a hierarchical process in which drug-target engagement is first encoded, then interpreted through transcriptomic state, and finally modulated by pharmacogenomic biomarkers. Transformer-based response models such as DeepTTA show how sequence-like drug and transcriptomic representations can be combined for prediction [9], while graph-based and pathway-guided methods demonstrate that biological structure can be retained inside neural networks [10, 11]. A hierarchical neural network could extend these ideas by explicitly nesting the drug, transcriptomic, and genomic modules in the same order that they influence cellular response. Such a model would be expected to improve interpretability by attributing resistance to the level at which the response signal becomes unfavorable.

Background

Principles of Drug Response in Cancer

Drug response in cancer depends on whether a compound engages relevant targets, whether the tumor cell is dependent on the affected pathway, and whether genomic alterations alter sensitivity or resistance. Target-centered modeling is important because kinase inhibition, receptor blockade, and pathway suppression can determine whether a drug has a plausible mechanism in a given tumor context [12]. Transcriptomic state adds another level, since pathway activity, lineage identity, and cell-state programs can influence whether target inhibition produces cytotoxic or cytostatic effects [13]. Pharmacogenomic alterations such as mutations and copy-number changes further shape response by modifying target dependency, bypass signaling, or intrinsic resistance, making them essential components of a patient-aware model [2, 3].

Public Pharmacogenomic Datasets

Public pharmacogenomic resources such as GDSC, CCLE, CTRP, and TCGA have enabled systematic modeling of drug sensitivity from molecular features and response labels. Deep learning surveys emphasize that these resources provide complementary views of pharmacogenomics, including gene expression, mutation, copy-number, and drug sensitivity measures, but also differ in assay conditions and biological coverage [1, 2]. Multi-omics response models such as MOLI use these data types to combine molecular layers for prediction while highlighting the need for careful integration across heterogeneous feature spaces [14]. Patient-linked resources can support translational evaluation, but models trained on cell lines should be interpreted cautiously because in vitro sensitivity is not equivalent to clinical response.

Deep Learning Approaches for Drug Response Prediction

Deep learning approaches for drug response prediction range from multilayer perceptrons on molecular profiles to graph neural networks that represent drugs or gene relationships explicitly. Models such as DeepDSC and convolutional response predictors learn nonlinear mappings between cancer cell line features and drug sensitivity, showing how neural networks can capture complex pharmacogenomic patterns [5, 15]. Graph convolutional models, including DeepCDR and later graph transformer variants, incorporate chemical structure, cell-line features, or biological networks to preserve some relational information during prediction [6, 16]. Attention-based multimodal encoders such as PaccMann further demonstrate that drug and transcriptomic representations can be fused while producing interpretable signals about molecular contributors to response [17, 18].

Hierarchical and Multi-Scale Neural Network Architectures

Hierarchical and multi-scale architectures are attractive for pharmacogenomic modeling because cancer drug response involves nested biological organization rather than independent feature blocks. DrugCell provides a clear example of a neural network structured by biological knowledge, where cellular subsystems constrain information flow and support mechanistic interpretation [7]. Pathway-guided and protein-association graph models similarly embed biological structure into response prediction, allowing pathway or network context to shape learned representations [10, 11]. Modular graph neural networks extend this idea by decomposing response prediction into biologically meaningful components that can support more transparent reasoning about drug action [19].

Interpretable Models for Biomarker Identification

Interpretability is essential in precision oncology because a predicted response is most useful when clinicians can understand which drug, pathway, or biomarker features support it. Multimodal attention models show how neural networks can highlight molecular features associated with anticancer compound sensitivity rather than only returning a response score [17]. Knowledge-guided architectures such as DRPreter use graph and transformer components to connect prediction with interpretable biological evidence [12]. Subcomponent-guided and explainable pathway-based models further illustrate how

attribution methods can identify candidate genes, pathways, or drug substructures that drive predicted sensitivity or resistance [20, 21].

Model Development Overview

High-Level Architecture

The proposed model consists of three interacting modules arranged in a biologically ordered hierarchy: a target-engagement module, a transcriptomic-context module, and a pharmacogenomic-context module. The target-engagement module encodes the drug's interaction with proteins or target classes, following the logic of drug-aware graph and chemical-structure encoders used in response prediction [6, 22]. Its latent output is passed into a transcriptomic module that contextualizes the drug signal using gene expression or pathway activity, an idea consistent with models that integrate compound and expression data for sensitivity estimation [13, 17]. The resulting representation is then modulated by a pharmacogenomic module that accounts for mutations and copy-number alterations before producing a personalized response prediction.

Figure 1 presents the proposed hierarchical neural network architecture that orders drug-target engagement, transcriptomic context, and pharmacogenomic modulation into an interpretable drug response prediction framework.

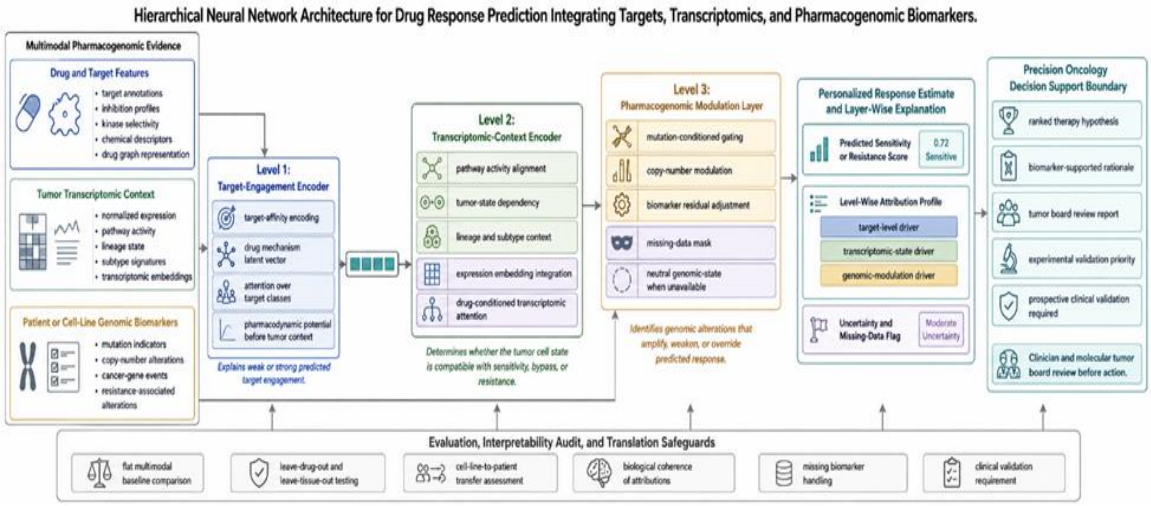

Figure 1. Hierarchical Neural Network Architecture for Drug Response Prediction Integrating Targets, Transcriptomics, and Pharmacogenomic Biomarkers.

Core Input Features

The core drug features include target annotations, biochemical inhibition profiles, kinase selectivity information, and molecular descriptors that summarize chemical or pharmacological properties. Drug-target and chemical representations can be encoded through graph neural networks, molecular encoders, or structured feature vectors, as illustrated by hybrid graph convolutional and graph transformer models for drug response prediction [6, 16]. Transcriptomic features may include normalized gene expression profiles, pathway activity scores, tumor subtype indicators, or embeddings learned from large-scale expression models such as scBERT, Geneformer, or scGPT [23-25]. Pharmacogenomic biomarkers can be represented as binary mutation indicators, copy-number scores, or gene-level alteration vectors, with the architecture designed to combine discrete genomic events with continuous expression signals [3, 14].

Design Principles

The model is designed around four principles: biological nesting, multimodal compatibility, missing-data tolerance, and interpretability. Biological nesting means that target engagement is not treated as equivalent to gene expression or mutation status, but instead serves as the mechanistic input that is progressively contextualized by tumor state and genomic background [7, 11]. Multimodal compatibility means the model can accept drug features, transcriptomic embeddings, and genomic biomarkers without requiring every modality to have the same scale or representation, reflecting lessons from multi-omics late-integration models [14]. Missing-data tolerance is important because pharmacogenomic biomarkers may be incomplete in some datasets, while interpretability ensures that the model can attribute predicted resistance to target, transcriptomic, or genomic levels rather than offering only a black-box score [12, 20].

Data Sources and Multi-Omics Feature Engineering

Cell-Line Pharmacogenomic Data

Cell-line pharmacogenomic datasets provide the primary development setting because they contain matched drug sensitivity and molecular profiles across many cancer models. Deep learning studies using GDSC, CCLE, and related panels show that gene expression, mutation, copy-number, and drug descriptors can be jointly modeled for sensitivity prediction [1, 2]. These resources support conceptual evaluation under cell-line-, drug-, and tissue-aware splits, but their limitations include batch

effects, assay differences, and incomplete representation of tumor microenvironment and patient pharmacology [8, 14]. Therefore, the proposed model should treat cell-line data as a structured preclinical training source rather than as a direct substitute for clinical response evidence.

Drug-Target Feature Extraction

For each drug, target information can be extracted as a vector of target affinities, inhibition strengths, selectivity scores, or binary target annotations. Graph and chemistry-aware response models show that representing the drug explicitly can improve the biological plausibility of predictions compared with using only cell-line molecular features [6, 22]. Target-aware models such as DRPreter and subcomponent-guided response predictors further suggest that the drug representation should support interpretation at the level of target classes, chemical substructures, or biological mechanisms [12, 20]. In the proposed hierarchy, these features form the lowest-level input because the drug's mechanism of action initiates the response cascade.

Pharmacogenomic Biomarker Encoding

Pharmacogenomic biomarkers can be encoded as mutation indicators, copy-number values, alteration burdens, or curated cancer-gene features that describe patient- or cell-line-specific genomic context. Integrated genomic-profile models and multi-omics late-integration approaches show that discrete genomic alterations can complement continuous transcriptomic data in predicting drug response [3, 14]. Copy-number alterations may be summarized at the gene level, while mutations may be represented as binary gene-level events or functionally weighted scores when variant interpretation is available. The encoding should preserve the distinction between genomic state and expression state because a resistance-conferring alteration can modulate drug response even when the transcriptomic profile appears permissive.

Hierarchical Neural Network Architecture

Drug-Target Engagement Module

The drug-target engagement module receives a structured drug representation and produces a latent vector describing the expected pharmacodynamic interaction between the compound and its molecular targets. Graph convolutional, graph transformer, and hybrid neural approaches provide useful design precedents for encoding compounds and drug-related relational features before response prediction [6, 16]. This module could use a feed-forward network for target-affinity vectors, a graph neural network for chemical structure, or an attention layer over target classes to emphasize the most relevant mechanisms [12]. Its output should not yet represent final sensitivity, but rather the drug's mechanistic potential before being interpreted in the tumor's cellular context.

Transcriptomic-Context Module

The transcriptomic-context module combines the target-engagement representation with gene expression, pathway activity, or transcriptomic embeddings to model whether the tumor cell state is compatible with drug sensitivity. Expression-based response models show that transcriptomic signatures can be highly informative for drug sensitivity, particularly when they capture pathway dependency or subtype-specific cellular programs [13]. Transformer-based and foundation-model embeddings such as DeepTTA, scBERT, Geneformer, and scGPT suggest that learned representations of expression data can provide richer context than individual genes alone [9, 23-25]. An attention gate could allow the model to weight transcriptomic programs according to the drug's target profile, so that downstream pathway effects are modeled after target engagement rather than independently from it.

Pharmacogenomic Modulation Module

The pharmacogenomic modulation module receives the transcriptomically contextualized drug signal and adjusts it using mutation and copy-number biomarkers. This design reflects the idea that genomic alterations can override, weaken, or amplify the response implied by drug-target engagement and expression state, as seen in integrated genomic and multi-omics response models [3, 14]. A gating mechanism, multiplicative interaction, or biomarker-conditioned residual layer could allow specific genomic events to modify the latent response trajectory rather than simply adding another flat feature block. Knowledge-guided and explainable response models indicate that this upper module should remain interpretable, enabling the model to identify when predicted resistance is associated with a specific mutation, amplification, deletion, or pathway-level genomic alteration [12, 21].

Table 1 defines how each biological level of the proposed hierarchy contributes a distinct modeling function, input structure, and interpretable response signal.

Table 1. Biological-Level Architecture of the Proposed Hierarchical Neural Network for Drug Response Prediction

Biological level	Model module	Primary input structure	Biological question encoded	Neural operation or design logic	Interpretable output contributed by the level
Drug mechanism level	Target-engagement encoder	Drug-target annotations, inhibition profiles, kinase selectivity, molecular	Does the compound plausibly engage a biologically relevant target mechanism?	Encodes target affinity, selectivity, chemical structure, or target-class attention into a	Indicates whether predicted nonresponse may arise from weak, nonspecific, or biologically

		descriptors, chemical graph features		pharmacodynamic latent vector	irrelevant target engagement
Cellular state level	Transcriptomic-context encoder	Gene expression, pathway activity scores, lineage markers, subtype signatures, transcriptomic embeddings	Is the tumor cell state dependent on, permissive to, or resistant to the drug mechanism?	Combines the target-engagement vector with expression-derived tumor context using attention, pathway alignment, or embedding fusion	Identifies whether sensitivity or resistance is driven by pathway activity, lineage state, bypass signaling, or unfavorable cellular programs
Genomic modulation level	Pharmacogenomic biomarker module	Mutation indicators, copy-number alterations, cancer-gene events, resistance biomarkers, alteration burden	Do patient- or cell-line-specific genomic alterations amplify, weaken, or override the contextualized response signal?	Applies biomarker-conditioned gating, multiplicative interaction, residual modulation, or mask-aware genomic adjustment	Explains when mutations, amplifications, deletions, or missing biomarker evidence modify the predicted response
Prediction level	Personalized response head	Hierarchically integrated latent representation from all three biological levels	What is the estimated likelihood or magnitude of drug sensitivity for this tumor-drug pair?	Produces continuous sensitivity, binary response probability, or ranked therapy hypothesis	Provides a response estimate that can be traced back to mechanistic, transcriptomic, and genomic drivers
Explanation level	Layer-wise attribution system	Module-specific latent vectors, attention weights, SHAP values, integrated gradients, uncertainty indicators	Which biological level most strongly supports or contradicts the predicted response?	Performs separate attribution within target, transcriptomic, and genomic modules rather than one global post hoc explanation	Distinguishes target-level, cell-state-level, and genomic-level rationales for sensitivity or resistance
Translation level	Decision-support reporting boundary	Prediction score, explanation profile, uncertainty flag, validation context	How should the model output be reviewed before clinical interpretation?	Converts model outputs into a tumor-board-facing interpretive report with validation and uncertainty safeguards	Supports ranked therapy hypotheses, biomarker review, experimental validation planning, and clinician oversight

Integrating Drug-Target, Transcriptomic, and Pharmacogenomic Modalities

Hierarchical Information Flow

The model is designed so that drug-target information is processed first, creating a pharmacodynamic latent vector that represents the compound's mechanistic potential before tumor context is considered. This design is consistent with response models that use compound structures, drug-target information, or chemical subcomponents as structured inputs rather than treating the drug as a simple identifier [20, 26]. The transcriptomic layer then contextualizes this drug signal by modeling whether the tumor's gene-expression state supports sensitivity, resistance, or pathway bypass [27, 28]. Finally, pharmacogenomic variants modulate the contextualized signal, allowing mutations and copy-number events to alter the expected response after target engagement and transcriptomic dependency have been represented.

Handling Missing Pharmacogenomic Data

Missing pharmacogenomic biomarkers are expected in many preclinical and clinical datasets, so the modulation module should be designed to operate even when mutation or copy-number features are incomplete. Multi-omics integration models such as MOLI show that separate modality encoders can support flexible integration when molecular data types differ across samples [14]. In the proposed hierarchy, missing genomic information could be represented by a learned neutral state, a mask-aware gating mechanism, or an uncertainty-aware embedding that prevents absent biomarkers from being interpreted as true negatives. This would allow the model to continue using drug-target and transcriptomic information while preserving the option to refine predictions when patient-specific genomic biomarkers become available.

Cross-Tissue and Tumor-Type Transferability

Cross-tissue transferability is important because a drug may share the same biochemical target across tumor types while the transcriptomic consequences of target inhibition differ by lineage. Transfer-learning approaches for cancer drug response suggest that model components trained on large pharmacogenomic screens can be adapted to patient or tumor-specific settings when the architecture separates general drug knowledge from context-specific biology [8, 29]. The proposed modular structure would allow the drug-target module to remain shared across tissues, while the transcriptomic-context module could be fine-tuned for tumor lineage, subtype, or pathway-state differences. Foundation transcriptomic models such as Geneformer and scGPT further support the idea that pretrained expression embeddings could provide reusable cellular representations for transfer across tumor contexts [24, 25].

Table 2 summarizes a modular design for improving cross-tissue transferability in cancer drug response models by separating shared pharmacological knowledge from tumor-specific transcriptional context.

Pharmacophore, 17(2) 2026, Pages 63-72

Table 2. Modular transfer-learning framework for cross-tissue drug response prediction

Component	Function	Learning Strategy	Role in Cross-Tissue Transfer	Example Approaches
Drug–target interaction module	Encodes drug properties and known molecular targets	Pretrained on large pharmacogenomic screening datasets	Shared across all tumor types to capture universal drug mechanism of action	DeepDTA-style encoders, graph neural networks, affinity predictors
Transcriptomic context module	Encodes tumor-specific gene expression profiles and pathway activity	Fine-tuned per lineage, subtype, or disease state	Adapts predictions to tissue-specific biology and signaling states	Gene expression encoders, pathway-aware neural layers
Cross-tissue adaptation layer	Aligns representations between training and target tumor distributions	Transfer learning / domain adaptation	Reduces distribution shift across cancer types	Adversarial domain adaptation, MMD-based alignment
Pretrained expression embeddings	Provides generalizable cellular representations	Self-supervised pretraining on large-scale transcriptomic datasets	Enables reuse of learned biological structure across tissues	Geneformer, scGPT
Prediction head	Integrates drug and tumor representations to predict response	Supervised learning on drug response labels	Produces final sensitivity/resistance output across cancers	Fully connected layers, multitask regression/classification

Model Interpretability and Biomarker Discovery

Layer-Wise Attribution

Layer-wise attribution would allow the model to explain whether a predicted resistant phenotype arises from weak target engagement, unfavorable transcriptomic state, or pharmacogenomic modulation. Attention-based compound sensitivity models demonstrate that neural networks can expose molecular features contributing to response rather than only producing a final prediction [17, 18]. Knowledge-guided and pathway-based models further show that attribution can be organized around biological units such as pathways, protein interactions, or drug subcomponents [11, 12]. In the proposed hierarchy, SHAP, integrated gradients, or attention weights should be applied separately to each module so that explanations remain aligned with the target, transcriptomic, and genomic levels of the architecture.

Biomarker Identification and Prioritization

Biomarker discovery would arise from repeated attribution analyses across drugs, cell lines, tumor types, and patient-derived profiles. Explainable models such as XGraphCDS and pathway-guided graph networks illustrate how response predictors can identify genes, pathways, and molecular mechanisms that plausibly contribute to sensitivity or resistance [21, 30]. Multimodal pathway-based models also suggest that combining omics differences with deep learning can prioritize candidate response features while preserving biological context [31]. The proposed model should treat such biomarkers as hypotheses for experimental validation, not as confirmed causal determinants, because attribution identifies predictive associations that require downstream biological testing.

Integration into Precision Oncology Decision Support

Translating Cell-Line Models to Patient Predictions

Translating a cell-line-trained model to patient prediction requires adaptation because clinical tumors differ from cultured cancer models in microenvironment, treatment history, clonal heterogeneity, and pharmacologic exposure. Few-shot and transfer-learning frameworks show that response predictors can be designed to move from high-throughput screens toward individual patient settings when patient-derived information is incorporated carefully [8]. Gene-expression-based inference models further support the use of transcriptomic profiles for drug sensitivity estimation, but their predictions should be interpreted as decision-support evidence rather than direct clinical proof [13]. The proposed hierarchy would therefore be evaluated conceptually as a translational framework in which drug-target knowledge remains stable, transcriptomic context is adapted, and pharmacogenomic biomarkers personalize the final prediction.

Supporting Molecular Tumor Boards

For molecular tumor boards, the model's value would depend not only on a sensitivity estimate but also on an explanation that can be reviewed by clinicians and molecular pathologists. Interpretable systems such as DRPreter and PaccMann show how response models can provide reasoning about drug mechanisms, expression features, or molecular contributors that underlie predictions [12, 18]. A hierarchical report could state that a therapy appears favorable because the drug strongly engages a relevant target, the tumor transcriptome indicates dependency on the affected pathway, and no resistance-associated genomic alteration is detected. Conversely, the model could flag predicted resistance when genomic modulation conflicts with otherwise favorable target and transcriptomic evidence.

Table 3 shows a structured interpretability framework that translates model outputs into clinically reviewable evidence streams supporting therapy selection or resistance identification in molecular tumor boards.

Table 3. Hierarchical interpretability framework for model-supported molecular tumor board decision-making

Evidence layer	Favorable signal	Resistance signal	Model interpretation	Clinical relevance
Drug–target interaction	Strong predicted target engagement or pathway inhibition	Weak binding affinity or off-target dominance	Pharmacological plausibility of response	Supports primary therapy selection
Transcriptomic context	Tumor shows pathway dependency or high target expression	Lack of pathway activation or alternative pathway dominance	Biological dependency alignment	Indicates likelihood of therapeutic sensitivity
Genomic alterations	Absence of known resistance mutations or amplifications	Presence of resistance-associated variants (e.g., pathway reactivation)	Genetic compatibility with drug mechanism	Refines patient stratification
Multi-omic concordance	Consistent signals across genomic, transcriptomic, and pharmacologic layers	Discordant signals across data modalities	Integrated confidence scoring	Strengthens or weakens overall recommendation
Resistance modeling	No predicted compensatory signaling	Predicted bypass pathway activation or feedback loops	Mechanistic resistance inference	Warns against likely non-response
Final interpretability output	Coherent multi-layer support for response	Conflicting evidence indicating reduced efficacy	Synthesized decision rationale	Guides MDT discussion and final review

Evaluation Strategy

Predictive Performance

Predictive performance should be evaluated against flat multimodal networks, regularized regression baselines, graph-based drug response models, and knowledge-guided architectures. Continuous response tasks would use correlation- and error-based metrics, while binary response tasks would use discrimination-based metrics, but the MDL framework should emphasize evaluation design rather than reporting experimental results. DeepCDR, DeepDSC, DrugCell, and graph convolutional response models provide appropriate conceptual comparators because they represent different ways of combining drug and molecular information [6, 7, 15, 22]. The key question is whether hierarchical ordering improves robustness and interpretability relative to architectures that combine features without explicitly modeling biological sequence.

Generalization across Cell Lines and Drugs

Generalization should be assessed under splits that test the model beyond familiar sample-drug combinations. Leave-cell-line-out and leave-tissue-out evaluation would test whether the transcriptomic and pharmacogenomic modules generalize to unseen tumor contexts, while leave-drug-out evaluation would test whether the target-engagement module supports prediction for therapies not observed during training [1, 16]. Transfer-oriented models such as TransCDR and few-shot drug response frameworks show why such evaluation is critical when moving from cell-line screens to patient-derived applications [8, 29]. The hierarchical model should therefore be judged by whether its modular structure supports generalization across both biological contexts and therapeutic mechanisms.

Interpretability Audit

An interpretability audit should examine whether the model's attributions are biologically coherent, stable across related samples, and consistent with known mechanisms of drug sensitivity or resistance. Pathway-guided, subcomponent-guided, and modular graph neural models provide useful precedents because they connect learned predictions to biological structures rather than relying solely on post hoc feature rankings [11, 19, 20]. The audit should test whether target-level explanations identify plausible drug mechanisms, transcriptomic explanations identify relevant pathway states, and genomic explanations identify mutations or copy-number changes that could modify response. This interpretability assessment should be treated as a qualitative and mechanistic evaluation rather than as proof that the model has discovered causal biomarkers.

Table 4 provides an evaluation framework for determining whether the hierarchical model offers advantages over flat multimodal prediction architectures in performance, generalization, interpretability, and clinical translation.

Table 4. Evaluation and Interpretability Framework for Testing Hierarchical Drug Response Prediction against Flat Multimodal Models

Evaluation domain	Core assessment question	Recommended comparison or audit	What the hierarchical model should demonstrate	Interpretation for precision oncology use
Predictive performance	Does biological ordering improve response prediction compared with flat feature concatenation?	Compare against flat multimodal neural networks, regularized regression, graph-based drug response models, and knowledge-guided architectures	Equal or improved correlation, error, discrimination, or calibration while preserving module-level interpretability	Strong performance alone is insufficient unless the prediction remains biologically explainable

Generalization across drugs	Can the model predict response for drugs or target mechanisms not seen during training?	Use leave-drug-out, leave-target-class-out, or chemically dissimilar split designs	The target-engagement module should transfer mechanistic information beyond memorized drug identifiers	Useful for therapy discovery only if predictions generalize across therapeutic mechanisms
Generalization across tumor contexts	Can the model adapt to unseen cell lines, tumor lineages, or patient-derived molecular profiles?	Use leave-cell-line-out, leave-tissue-out, tumor-subtype holdout, and transfer-learning evaluations	The transcriptomic-context module should retain tumor-state reasoning across different biological backgrounds	Supports translational relevance when the model is not limited to familiar cell-line distributions
Biomarker modulation validity	Do mutations and copy-number events modify predictions in biologically plausible ways?	Perform counterfactual biomarker perturbation, known resistance-event testing, and alteration-specific subgroup review	Genomic modulation should change predictions when biomarkers plausibly alter target dependency, bypass signaling, or resistance	Prevents mutation and copy-number features from becoming uninterpretable flat covariates
Missing-data robustness	Does the model distinguish unavailable biomarker data from true negative biomarker evidence?	Compare complete-case, mask-aware, neutral-state, and uncertainty-aware inference settings	Prediction should remain stable when genomic data are missing while flagging uncertainty appropriately	Essential for clinical settings where molecular profiling depth varies across patients
Explanation coherence	Are layer-wise attributions biologically stable and clinically interpretable?	Audit target-level, transcriptomic-level, and genomic-level explanations separately using known mechanisms and pathway knowledge	Explanations should identify plausible drug targets, expression programs, and genomic modifiers without collapsing into one opaque ranking	Enables molecular tumor boards to review why a therapy appears favorable or unfavorable
Cell-line-to-patient translation	Does the model remain useful when moving from pharmacogenomic screens to patient decision support?	Evaluate external patient-derived datasets, organoid models, retrospective clinical cohorts, or few-shot adaptation settings	The model should preserve mechanistic structure while acknowledging assay, microenvironment, and treatment-history differences	Positions the model as decision-support evidence rather than direct clinical proof
Clinical governance	Are predictions appropriately bounded by validation, uncertainty, and human review?	Review calibration, uncertainty reporting, prospective validation requirements, and tumor-board usability	Outputs should include response estimate, explanation profile, confidence context, and validation limitations	Supports responsible precision oncology use without presenting the model as an autonomous treatment selector

Limitations

Cell-Line to Patient Translation Gap

The most important limitation is that cell-line pharmacogenomic screens cannot fully represent clinical tumors. Even sophisticated deep learning models trained on GDSC, CCLE, or related resources may learn associations shaped by culture conditions, assay platforms, and missing microenvironmental features rather than treatment response in patients [1, 2]. Patient-focused adaptation strategies can help, but clinical labels are often heterogeneous because they depend on prior therapy, combination regimens, dosing, immune context, and disease stage [8]. Therefore, the hierarchical model should be regarded as a translational decision-support architecture that requires prospective validation before clinical use.

Static Representation of Drug Targets

A second limitation is that drug-target features are often static summaries of pharmacology, whereas real tumors exhibit dynamic feedback, adaptive pathway rewiring, and context-dependent target engagement. Chemical-structure and target-aware models can encode drug mechanism more richly than simple identifiers, but they may still miss time-dependent compensatory signaling and microenvironmental modulation [26, 32]. Transcriptomic foundation models and pathway embeddings may partly capture cellular state, yet they usually represent snapshots rather than longitudinal drug perturbation trajectories [23, 25]. Future versions of the hierarchy should therefore incorporate perturbational, temporal, and in vivo data to model how drug response evolves after treatment exposure.

Conclusion

A hierarchical neural network for drug response prediction would model anticancer sensitivity as a nested biological process. In this framework, drug-target engagement forms the mechanistic base, transcriptomic context determines whether that mechanism is relevant in the tumor cell state, and pharmacogenomic biomarkers modulate the final patient-specific response estimate.

The main strength of this design is that it aligns the neural architecture with the biological organization of precision oncology. Rather than flattening all inputs into a single vector, the model separates target engagement, cellular context, and genomic modulation so that each level can contribute both to prediction and interpretation.

The remaining challenges are substantial. Cell-line-to-patient translation, incomplete biomarker data, dynamic pharmacology, and clinical validation all limit the immediate use of such models for treatment selection.

Future work should incorporate temporal perturbation profiles, patient-derived tumor models, and in vivo response data to make the hierarchy more clinically realistic. Community benchmarking against flat multimodal baselines would help determine whether biologically structured architectures offer practical advantages for precision oncology decision support.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Baptista D, Ferreira PG, Rocha M. Deep learning for drug response prediction in cancer. Brief Bioinform. 2021;22(1):360-79.
2. Chiu YC, Chen HI, Gorthi A, Mostavi M, Zheng S, Huang Y, et al. Deep learning of pharmacogenomics resources: moving towards precision oncology. Brief Bioinform. 2020;21(6):2066-83.
3. Chiu YC, Chen HI, Zhang T, Zhang S, Gorthi A, Wang LJ, et al. Predicting drug response of tumors from integrated genomic profiles by deep neural networks. BMC Med Genomics. 2019;12(Suppl 1):18.
4. Sakellaropoulos T, Vougas K, Narang S, Koinis F, Kotsinas A, Polyzos A, et al. A deep learning framework for predicting response to therapy in cancer. Cell Rep. 2019;29(11):3367-73.
5. Liu P, Li H, Li S, Leung KS. Improving prediction of phenotypic drug response on cancer cell lines using deep convolutional network. BMC Bioinformatics. 2019;20(1):408.
6. Liu Q, Hu Z, Jiang R, Zhou M. DeepCDR: a hybrid graph convolutional network for predicting cancer drug response. Bioinformatics. 2020;36(Suppl_2):i911-8.
7. Kuenzi BM, Park J, Fong SH, Sanchez KS, Lee J, Kreisberg JF, et al. Predicting drug response and synergy using a deep learning model of human cancer cells. Cancer Cell. 2020;38(5):672-84.
8. Ma J, Fong SH, Luo Y, Bakkenist CJ, Shen JP, Mourragui S, et al. Few-shot learning creates predictive models of drug response that translate from high-throughput screens to individual patients. Nat Cancer. 2021;2(2):233-44.
9. Jiang L, Jiang C, Yu X, Fu R, Jin S, Liu X. DeepTTA: a transformer-based model for predicting cancer drug response. Brief Bioinform. 2022;23(3):bbac100.
10. Zhu Y, Ouyang Z, Chen W, Feng R, Chen DZ, Cao J, et al. TGSA: protein–protein association-based twin graph neural networks for drug response prediction with similarity augmentation. Bioinformatics. 2022;38(2):461-8.
11. Deng L, Cai Y, Zhang W, Yang W, Gao B, Liu H. Pathway-guided deep neural network toward interpretable and predictive modeling of drug sensitivity. J Chem Inf Model. 2020;60(10):4497-505.
12. Shin J, Piao Y, Bang D, Kim S, Jo K. DRPreter: interpretable anticancer drug response prediction using knowledge-guided graph neural networks and transformer. Int J Mol Sci. 2022;23(22):13919.
13. Chawla S, Rockstroh A, Lehman M, Ratther E, Jain A, Anand A, et al. Gene expression based inference of cancer drug sensitivity. Nat Commun. 2022;13(1):5680.
14. Sharifi-Noghabi H, Zolotareva O, Collins CC, Ester M. MOLI: multi-omics late integration with deep neural networks for drug response prediction. Bioinformatics. 2019;35(14):i501-9.
15. Li M, Wang Y, Zheng R, Shi X, Li Y, Wu FX, et al. DeepDSC: a deep learning method to predict drug sensitivity of cancer cell lines. IEEE/ACM Trans Comput Biol Bioinform. 2021;18(2):575-82.
16. Chu T, Nguyen TT, Hai BD, Nguyen QH, Nguyen T. Graph transformer for drug response prediction. IEEE/ACM Trans Comput Biol Bioinform. 2023;20(2):1065-72.
17. Manica M, Oskooei A, Born J, Subramanian V, Sáez-Rodríguez J, Rodriguez Martinez M. Toward explainable anticancer compound sensitivity prediction via multimodal attention-based convolutional encoders. Mol Pharm. 2019;16(12):4797-806.
18. Cadow J, Born J, Manica M, Oskooei A, Rodríguez Martínez M. PaccMann: a web service for interpretable anticancer compound sensitivity prediction. Nucleic Acids Res. 2020;48(W1):W502-8.
19. Campana PA, Prasse P, Lienhard M, Thedinga K, Herwig R, Scheffer T. Cancer drug sensitivity estimation using modular deep Graph Neural Networks. NAR Genom Bioinform. 2024;6(2):lqae043.
20. Liu X, Zhang W. A subcomponent-guided deep learning method for interpretable cancer drug response prediction. PLoS Comput Biol. 2023;19(8):e1011382.
21. Wang Y, Yu X, Gu Y, Li W, Zhu K, Chen L, et al. XGraphCDS: an explainable deep learning model for predicting drug sensitivity from gene pathways and chemical structures. Comput Biol Med. 2024;168:107746.

22. Nguyen T, Nguyen GT, Nguyen T, Le DH. Graph convolutional networks for drug response prediction. IEEE/ACM Trans Comput Biol Bioinform. 2022;19(1):146-54.
23. Yang F, Wang W, Wang F, Fang Y, Tang D, Huang J, et al. scBERT as a large-scale pretrained deep language model for cell type annotation of single-cell RNA-seq data. Nat Mach Intell. 2022;4(10):852-66.
24. Theodoris CV, Xiao L, Chopra A, Chaffin MD, Al Sayed ZR, Hill MC, et al. Transfer learning enables predictions in network biology. Nature. 2023;618(7965):616-24.
25. Cui H, Wang C, Maan H, Pang K, Luo F, Duan N, et al. scGPT: toward building a foundation model for single-cell multi-omics using generative AI. Nat Methods. 2024;21(8):1470-80.
26. Zuo Z, Wang P, Chen X, Tian L, Ge H, Qian D. SWnet: a deep learning model for drug response prediction from cancer genomic signatures and compound chemical structures. BMC Bioinformatics. 2021;22(1):434.
27. Kim S, Bae S, Piao Y, Jo K. Graph convolutional network for drug response prediction using gene expression data. Mathematics. 2021;9(7):772.
28. Li P, Jiang Z, Liu T, Liu X, Qiao H, Yao X. Improving drug response prediction via integrating gene relationships with deep learning. Brief Bioinform. 2024;25(3):bbae153.
29. Xia X, Zhu C, Zhong F, Liu L. TransCDR: a deep learning model for enhancing the generalizability of drug activity prediction through transfer learning and multimodal data fusion. BMC Biol. 2024;22(1):227.
30. Yoon HJ, Lee M. GCNPath: introspecting drug response prediction with pathway-guided graph convolution networks. Commun Biol. 2026;9(1):224.
31. Wu Y, Chen M, Qin Y. Anticancer drug response prediction integrating multi-omics pathway-based difference features and multiple deep learning techniques. PLoS Comput Biol. 2025;21(3):e1012905.
32. Chen Y, Zhang L. Hi-GeoMVP: a hierarchical geometry-enhanced deep learning model for drug response prediction. Bioinformatics. 2024;40(4):btae204.